

SCSI to SATA RAID Subsystem

DVA-08E | DVA-08K | DVA-16K



BOSCH

en Administrator's Manual

Table of Contents

1	RAID Functions: Introduction	9
1.1	Logical Drive	9
1.2	Logical Volume	9
1.3	RAID Levels	10
1.3.1	NRAID	11
1.3.2	JBOD	11
1.3.3	RAID0	12
1.3.4	RAID1	12
1.3.5	RAID (0+1)	13
1.3.6	RAID3	14
1.3.7	RAID5	14
1.4	Spare Drives	15
1.4.1	Local Spare Drives	15
1.4.2	Global Spare Drives	15
1.5	Rebuild	17
1.5.1	Automatic Rebuild	17
1.5.2	Manual Rebuild	18
1.5.3	Concurrent Rebuild in RAID (0+1)	18
1.6	Logical Volume (Multi-Level RAID)	19
1.6.1	What Is a Logical Volume?	19
1.6.2	Can Spare Drives Assigned to a Logical Volume?	21
1.6.3	Limitations	21
1.6.4	Partitioning the Logical Drive or Partitioning the Logical Volume?	21
1.6.5	RAID Expansion with Logical Volume?	22
1.6.6	A Logical Volume with Logical Drives of Different Levels?	22
2	RAID Planning	23
2.1	Considerations	23
2.1.1	How many physical drives do you have?	23
2.1.2	Limitations? Does my OS or HBA support 64-bit LBA?	23
2.1.3	How many drives on each drive channel?	23
2.1.4	How many drives would you like to appear to the host computer?	23
2.1.5	What kind of host application?	23
2.1.6	Multi-path, hub, or switch?	24
2.1.7	Optimization mode	24
2.1.8	What RAID level?	24
2.1.9	Any spare drives?	24
2.2	Configuring the Array	25
2.2.1	Starting a RAID System	25
2.3	Operation Theory	26
2.3.1	I/O Channel, Host ID, and LUN	26
2.3.2	Grouping Drives into an Array	26
2.3.3	Making Arrays Available to Hosts	28
2.4	Tunable Parameters	29
3	Accessing the Array: Serial Port, Ethernet, and Access Security	33
3.1	RS-232C Serial Port	33
3.1.1	Configuring RS-232C Connection via LCD Keypad Panel	33
3.1.2	Starting RS-232C Terminal Emulation	34
3.2	Communication Parameters	35
3.2.1	RS-232 Port Configuration	36
3.2.2	Terminal Emulation:	36
3.2.3	Baud Rate	36
3.2.4	Internet Protocol <TCP/IP>	37
3.2.5	View Statistics	37
3.2.6	Set IP Address	37
3.3	Out-of-Band via Ethernet	38
3.3.1	What is the "Disk Reserved Space?"	39
3.3.2	Other Concerns	39
3.3.3	Connecting Ethernet Port:	39

3.3.4	Configuring the Controller	39
3.4	Telnet Connection	40
3.5	Secure Link over SSH	41
4	LCD Screen Messages	42
4.1	The Initial Screen	42
4.2	Quick Installation Screen	42
4.3	Logical Drive Status	43
4.4	Logical Volume Status	44
4.5	Drive Status	44
4.6 Channel Status	45
4.7	Controller Voltage and Temperature	46
4.8	View and Edit Event Logs	47
5	LCD Keypad Operation	48
5.1	Power On the RAID Enclosure	48
5.2	Caching Parameters	48
5.2.1	Optimization Modes	48
5.2.2	Optimization Mode and Stripe Size	49
5.2.3	Optimization for Random or Sequential I/Os	49
5.2.4	Write-back/Write-through Cache Enable/Disable	49
5.2.5	Periodic Cache Flush	50
5.3	View Connected Drives:	51
5.4	Creating a Logical Drive	52
5.4.1	Choosing a RAID Level	52
5.4.2	Choosing Member Drives	52
5.4.3	Logical Drive Preferences	52
5.4.4	Maximum Drive Capacity	53
5.4.5	Spare Drive Assignments	53
5.4.6	Disk Reserved Space	53
5.4.7	Write Policy	53
5.4.8	Initialization Mode	54
5.4.9	Stripe Size	54
5.4.10	Beginning Initialization	54
5.5	Creating a Logical Volume	56
5.5.1	Initialization Mode	56
5.5.2	Write Policy	56
5.6	Partitioning a Logical Drive/Logical Volume	58
5.7	Mapping a Logical Volume/Logical Drive to Host LUN	59
5.8	Assigning Spare Drive and Rebuild Settings	61
5.8.1	Adding a Local Spare Drive	61
5.8.2	Adding a Global Spare Drive	61
5.8.3	Rebuild Settings	62
5.9	Viewing and Editing Logical Drives and Drive Members	63
5.9.1	Deleting a Logical Drive	63
5.9.2	Deleting a Partition of a Logical Drive	64
5.9.3	Assigning a Name to a Logical Drive	65
5.9.4	Rebuilding a Logical Drive (Manual Rebuild)	65
5.9.5	Regenerating Logical Drive Parity	66
5.9.6	Media Scan	67
5.9.7	Write Policy	68
5.10	Viewing and Editing Host LUNs	69
5.11	Viewing and Editing Drives	69
5.12	Viewing and Editing Channels	70
5.12.1	Viewing IDs	70
5.12.2	Adding a Channel ID	70
5.12.3	Deleting a Channel ID	71
5.12.4	Data Rate	72
5.13	System Functions	74
5.13.1	Mute Beeper	74
5.13.2	Password	74
5.13.3	Reset Controller	75

5.13.4	Shutdown Controller	75
5.13.5	Controller Maintenance	75
5.13.6	Saving NVRAM to Disks	76
5.13.7	Restore NVRAM from Disks	76
5.13.8	Adjust LCD Contrast	77
5.14	Controller Parameters	78
5.14.1	Controller Name	78
5.14.2	LCD Title Display Controller Name	78
5.14.3	Password Validation Timeout	78
5.14.4	Controller Unique Identifier	78
5.14.5	Controller Date and Time	79
5.15	Drive Utilities	80
5.15.1	Drive Read/Write Test	80
6	Terminal Screen Messages	82
6.1	The Initial Screen	82
6.2	Main Menu	83
6.3	Quick Installation	83
6.4	Logical Drive Status	84
6.5	Logical Volume Status	85
6.6	Drive Status	86
6.7	Channel's Status	87
6.8	Controller Voltage and Temperature	88
6.9	Viewing Event Logs on the Screen	89
7	Terminal Operation	90
7.1	Power on RAID Enclosure	90
7.2	Caching Parameters	91
7.2.1	Optimization Modes	91
7.2.2	Optimization Mode and Stripe Size	91
7.2.3	Optimization for Random or Sequential I/Os	92
7.2.4	Write-Back Cache Enable/Disable	92
7.2.5	Periodic Cache Flush	93
7.3	Viewing the Connected Drives	94
7.4	Creating a Logical Drive	95
7.4.1	Choosing a RAID Level	95
7.4.2	Choosing Member Drives	95
7.4.3	Logical Drive Preferences	96
7.5	Creating a Logical Volume	99
7.6	Partitioning a Logical Drive/Logical Volume	101
7.7	Mapping a Logical Volume to Host LUNs	103
7.8	Assigning a Spare Drive and Rebuild Settings	105
7.8.1	Adding Local Spare Drive	105
7.8.2	Adding a Global Spare Drive	106
7.9	Viewing and Editing Logical Drives and Drive Members	107
7.9.1	Deleting a Logical Drive	108
7.9.2	Deleting a Partition of a Logical Drive	108
7.9.3	Assigning a Name to a Logical Drive	109
7.9.4	Rebuilding a Logical Drive	109
7.9.5	Expand Logical Drive	110
7.9.6	Add Drive	110
7.9.7	Regenerating Logical Drive Parity	110
7.9.8	Media Scan	111
7.9.9	Write Policy	111
7.10	Viewing and Editing Host LUNs	112
7.11	Viewing and Editing Drives	113
7.11.1	Deleting Spare Drive (Global / Local Spare Drive)	114
7.11.2	Disk Reserved Space	114
7.11.3	Identifying Drives	115
7.11.4	Media Error Reassign Count	117
7.12	Viewing and Editing Channels	118
7.12.1	Viewing and Editing IDs - Host Channel	118

7.12.2	Adding an ID	119
7.12.3	Deleting an ID	119
7.12.4	Data Rate (Channel Bus)	120
7.12.5	View Chip Information	120
7.13	System Functions	121
7.13.1	Mute Beeper	121
7.13.2	Change Password	121
7.13.3	Changing the Password	122
7.13.4	Setting a New Password	122
7.13.5	Disabling the Password	122
7.13.6	Reset Controller	123
7.13.7	Shutdown Controller	123
7.13.8	Adjust LCD Contrast	123
7.14	Controller Parameters	125
7.14.1	Controller Name	125
7.14.2	LCD Title Display - Controller Name	125
7.14.3	Saving NVRAM to Disks	126
7.14.4	Restore NVRAM from Disks	126
7.14.5	Password Validation Timeout	127
7.14.6	Controller Unique Identifier	127
7.14.7	Set Controller Date and Time	128
7.15	View Drive Information	129
7.16	Drive Utilities	129
7.16.1	Drive Read/Write Test	129
8	Host-side and Drive-side Parameters	131
8.1	Host-side Parameters	131
8.1.1	Maximum Concurrent Host LUN Connection ("Nexus" in SCSI)	131
8.1.2	Number of Tags Reserved for Each Host-LUN Connection	132
8.1.3	Maximum Queued I/O Count	133
8.1.4	LUNs per Host ID	133
8.1.5	LUN Applicability	134
8.1.6	Peripheral Device Type	134
8.1.7	In-band (SCSI channel)	134
8.1.8	Peripheral Device Type Parameters for Various Operating Systems	135
8.1.9	Cylinder/Head/Sector Mapping	137
8.2	Drive-side Parameters	139
8.2.1	Disk Access Delay Time	139
8.2.2	Drive I/O Timeout	140
8.2.3	Maximum Tag Count: Tag Command Queuing (TCQ) and Native Command Queuing (NCQ) Support	141
8.2.4	Drive Spindown Idle Delay Period	142
8.2.5	Drive Delayed Write	142
9	Enclosure Management	143
9.1	Enclosure Monitoring	143
9.1.1	Enclosure Devices	143
9.1.2	Auto Shutdown: Elevated Temperature	145
9.1.3	RAID Controller Self-monitoring	147
9.2	SAF-TE and SES Enclosure Monitoring	149
9.2.1	Periodic SAF-TE and SES Device Check Time	149
9.2.2	Enclosure Devices Status: View Peripheral Device Status	149
10	Data Integrity	151
10.1	Failed Drive Detection	151
10.1.1	Periodic Drive Check Time	151
10.1.2	Periodic Auto-Detect Failure Drive Swap Check Time	152
10.1.3	Auto-Assign Global Spare Drive	152
10.2	Scheduled Maintenance	154
10.2.1	Task Scheduler	154
10.2.2	Task Scheduler Options	155
10.2.3	Confirming the Creation of a Task Schedule	157
10.3	Regenerating Logical Drive Parity	158
10.3.1	Overwrite Inconsistent Parity	158

10.3.2	Generate Check Parity Error Event	158
10.4	Disk Array Parameters	159
10.4.1	Rebuild Priority	159
10.4.2	Verification on Writes	159
11	Array Expansion.	161
11.1	Overview	161
11.2	Notes on Expansion.	161
11.2.1	Added Capacity	161
11.2.2	Size of the New Drive	161
11.2.3	Applicable Arrays.	161
11.2.4	Interruption to the Process.	162
11.3	Mode 1 Expansion:	163
11.4	Mode 2 Expansion.	166
11.5	Making Use of the Added Capacity: Expand Logical Drive	168
11.6	Expand Logical Volume	170
11.7	Configuration Example: Volume Extension in Windows 2000®	170
12	S.M.A.R.T. Configuration	176
12.1	Overview	176
12.2	Clone Failing Drive	176
12.3	S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology)	179
12.4	Configuration Procedure.	181
13	Implementations for AV Applications	184
13.1	Maximum Drive Response Time	184
13.2	AV Optimization Mode.	185
14	Firmware Functionality Specifications	186
14.1	Basic RAID Management	186
14.2	Advanced Features	189
14.3	Caching Operation	191
14.4	RAID Expansion	192
14.5	S.M.A.R.T. Support	192
14.6	Redundant Controller	193
14.7	Data Safety	194
14.8	System Security.	195
14.9	Environment Management.	195
14.10	User Interface	196
14.11	High Availability.	196
15	System Functions - Upgrading Firmware.	197
15.1	Upgrading Firmware	197
15.1.1	Background RS-232C Firmware Download: Single Controller	197
15.1.2	Note for Redundant Controller Firmware Upgrade:	197
15.2	Upgrading Firmware Using RAIDWatch Manager.	197
15.2.1	Establish the In-band SCSI connection in RAIDWatch Manager	197
15.2.2	Upgrade Both Boot Record and Firmware Binaries.	197
15.2.3	Upgrade the Firmware Binary Only	198
15.3	Upgrading Firmware Using RS-232C Terminal Emulation	198
15.3.1	Establishing the connection for the RS-232C Terminal Emulation	198
15.3.2	Upgrading Both Boot Record and Firmware Binaries	198
15.3.3	Upgrading the Firmware Binary Only	199
16	Event Messages.	200
16.1	Event Index	200
16.1.1	Logical Drive Events	200
16.1.2	Channel and Individual Drive Events.	202
16.1.3	General Target Events	204
16.1.4	Controller Events.	205
16.2	Event Description	206
16.2.1	Logical Drive Events	206
16.2.2	Channel and Individual Drive Events	215
16.2.3	General Target Events.	223


16.2.4 Controller Events 228

1 RAID Functions: Introduction

Redundant Arrays of Independent Disks, or RAID, offers the following advantages: availability, capacity, and performance. Choosing the right RAID level and drive failure management can increase capacity and performance, subsequently increasing availability. Bosch's external RAID controllers and subsystems provide complete RAID functionality and enhanced drive failure management.

1.1 Logical Drive

The advantages mentioned above are achieved by creating “logical drives.” A logical drive is an array of independent physical drives. The logical drive appears to the host as a contiguous volume, the same as a local hard disk drive does.



Logical Drive

Fig. 1.1 Logical Drive

The following section describes the different methods to create logical arrays of disk drives, such as spanning, mirroring and data parity. These methods are referred to as “RAID levels.”

1.2 Logical Volume

The concept of a logical volume is very similar to that of a logical drive. A logical volume is the combination of one or several logical drives. These logical drives are combined into a larger capacity using the RAID0 method (striping). When data is written to a logical volume, it is first broken into data segments and then striped across different logical drives in a logical volume. Each logical drive then distributes data segments to its member drives according to the specific RAID level it is composed of.


The member logical drives can be composed of the same RAID level or each of a different RAID level. A logical volume can be divided into a maximum of 8 partitions. During operation, a host computer sees a non-partitioned logical volume or a partition of a logical volume as one single physical drive.

1.3 RAID Levels

- Using a RAID storage subsystem has the following advantages:
- Provides disk spanning by weaving all connected drives into one single volume.
 - Increases disk access speed by breaking data into several blocks when reading/writing to several drives in parallel. With RAID, storage speed increases as more drives are added as the channel bus allows.
 - Provides fault-tolerance by mirroring or distributing parity across disk drives.

RAID Level	Description	Capacity	Data Availability
NRAID	Non-RAID	N	None
RAID0	Disk Striping	N	Less than one single drive
RAID1 (0+1)	Mirroring Plus Striping (if N>1)	N/2	high ==RAID5
RAID3	Striping with Parity on dedicated disk	N-1	high ==RAID5
RAID5	Striping with interspersed parity	N-1	high ==RAID5
RAID 10 (Logical Volume)	Striping with RAID1 logical drives	/	high >>RAID5
RAID 30 (Logical Volume)	Striping with RAID3 logical drives	/	high >>RAID5
RAID 50 (Logical Volume)	Striping with RAID5 logical drives	/	high >>RAID5

Table 1.1 A Brief on RAID Levels



NOTICE!

Drives on different channels can be included in a logical drive, and logical drives of different RAID levels can be used to compose a logical volume. There are more combinations than RAID 10, 30, and 50.

RAID Level	Performance Sequential	Performance Random	
NRAID	Drive	Drive	
RAID0	R: Highest W: Highest	R: High W: Highest	
RAID1 (0+1)	R: High W: Medium	R: Medium W: Low	
RAID3	R: High W: Medium	R: Medium W: Low	
RAID5	R: High W: Medium	R: High W: Low	

1.3.1

NRAID

Disk Spanning

NRAID stands for Non-RAID. The capacity of all drives is combined to become one logical drive (no block striping). In other words, the capacity of the logical drive is the total capacity of the physical member drives. NRAID does not provide data redundancy.

NRAID	
Minimum Disks Required	1
Capacity	N
Redundancy	No

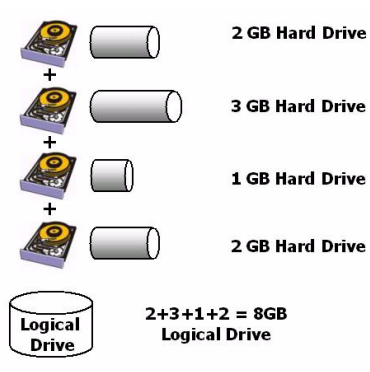


Fig. 1.2 NRAID

1.3.2

JBOD

Single Drive Control

JBOD stands for Just a Bunch of Drives. The controller treats each drive as a stand-alone disk; therefore, each drive is an independent logical drive. JBOD does not provide data redundancy.

JBOD	
Minimum Disks Required	1
Capacity	1
Redundancy	No




Fig. 1.3 JBOD

1.3.3

RAID0

Disk Striping

RAID0 provides the highest performance but no redundancy. Data in the logical drive is striped (distributed) across several physical drives.

RAID0	
Minimum Disks Required	2
Capacity	N
Redundancy	No




Fig. 1.4 RAID0

1.3.4

RAID1

Disk Mirroring

RAID1 mirrors the data stored in one hard drive to another. RAID1 can only be performed with two hard drives. If there are more than two hard drives, RAID (0+1) will be automatically applied.

RAID1	
Minimum Disks Required	2
Capacity	N/2
Redundancy	Yes




Fig. 1.5 RAID1

1.3.5

RAID (0+1)**Disk Striping with Mirroring**

RAID (0+1) combines RAID0 and RAID1 - Mirroring and Striping. RAID (0+1) allows multiple drive failure because of the full redundancy of the hard drives. If there are more than two hard drives assigned to perform RAID1, RAID (0+1) will be automatically applied.

RAID (0+1)	
Minimum Disks Required	4
Capacity	N/2
Redundancy	Yes




Fig. 1.6 RAID (0+1)

**NOTICE!**

“RAID (0+1)” will not appear in the list of RAID levels supported by the controller. If you wish to perform RAID1, the subsystem will determine whether to perform RAID1 or RAID (0+1). This will depend on the number of drives that have been selected as the members of the logical drive.

1.3.6

RAID3

Disk Striping with Dedicated Parity Disk

RAID3 performs Block Striping with Dedicated Parity. One drive member is dedicated to storing the parity data. When a drive member fails, the controller can recover/regenerate the lost data of the failed drive by comparing and re-calculating data on the remaining drives.

RAID (0+1)	
Minimum Disks Required	3
Capacity	N-1
Redundancy	Yes




Fig. 1.7 FRAID3

1.3.7

RAID5

Striping with Interspersed Parity

RAID5 is similar to RAID3 but the parity data is not stored in a dedicated hard drive. Parity information is interspersed across the drive array. In the event of a drive failure, the controller can recover/regenerate the lost data of the failed drive by comparing and re-calculating data on the remaining drives.

RAID5	
Minimum Disks Required	3
Capacity	N-1
Redundancy	Yes

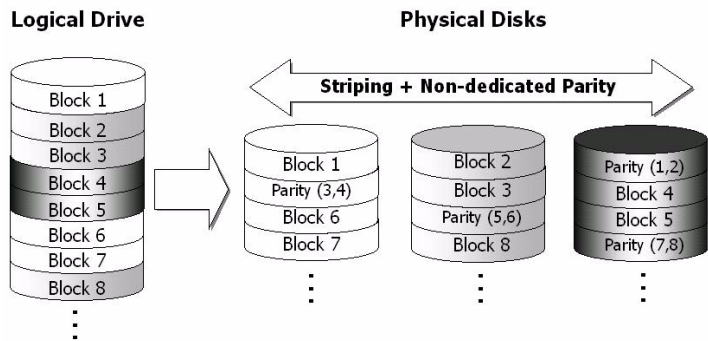


Fig. 1.8 RAID5

1.4 Spare Drives

1.4.1 Local Spare Drives

A Local Spare Drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the Local Spare Drive becomes a member drive and automatically starts to rebuild.

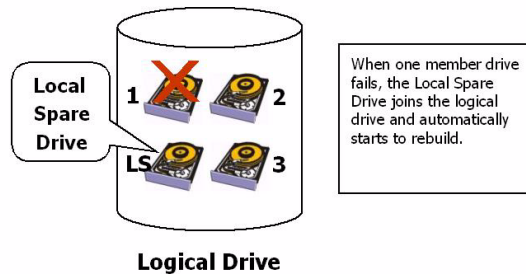


Fig. 1.9 Local (Dedicated) Spare

1.4.2 Global Spare Drives

A Global Spare Drive serves more than one specified logical drive. When a member drive from any of the logical drives fails, the Global Spare Drive will join that logical drive and automatically start to rebuild.

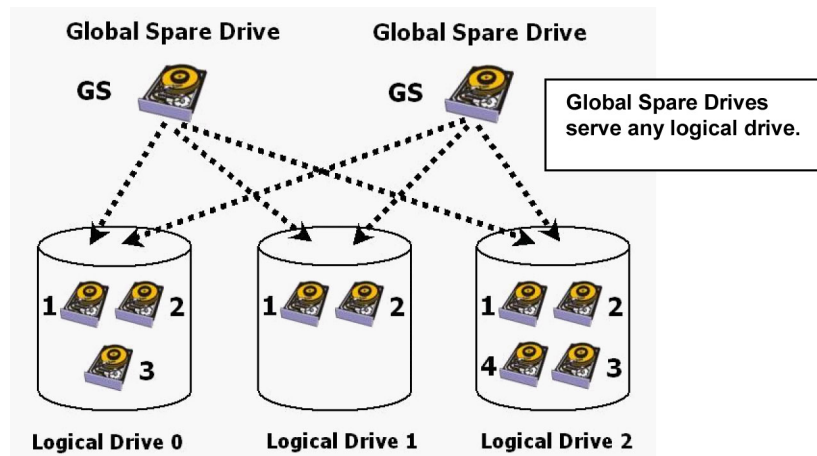


Fig. 1.10 Global Spare

The example on the left provides both Local Spare Drive and Global Spare Drive functions. On certain occasions, applying these two functions together will better fit various needs. Take note that the Local Spare Drive always has higher priority than the Global Spare Drive.

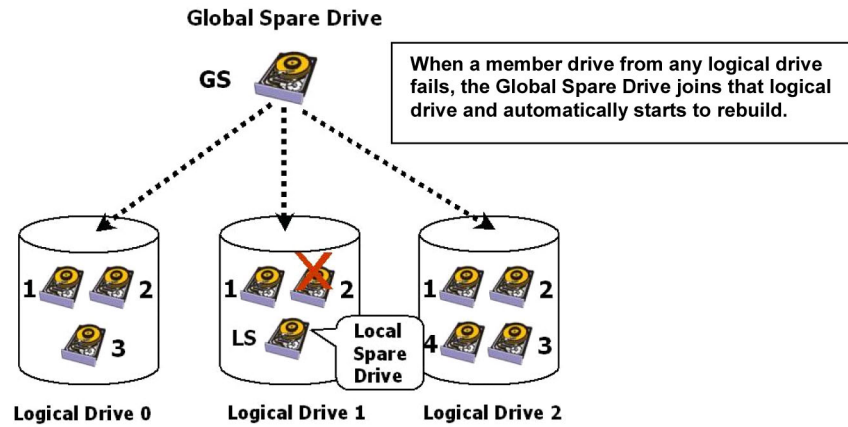


Fig. 1.11 Global Spare Rebuild

In the example shown below, the members of Logical Drive 0 are 9GB drives, and the members in Logical Drives 1 and 2 are 4GB drives.

It is not possible for the 4GB Global Spare Drive to join Logical Drive 0 because of its insufficient capacity. However, using a 9GB drive as the Global Spare Drive for a failed drive that comes from Logical Drive 1 or 2 will bring a huge amount of excess capacity since these logical drives require 4GB only. As shown on the left, the 9GB Local Spare Drive will aid Logical Drive 0 once a drive in this logical drive fails. If the failed drive is in Logical Drive 1 or 2, the 4GB Global Spare drive will immediately give aid to the failed drive.




Fig. 1.12 Mixing Local and Global Spares

1.5

Rebuild

1.5.1

Automatic Rebuild

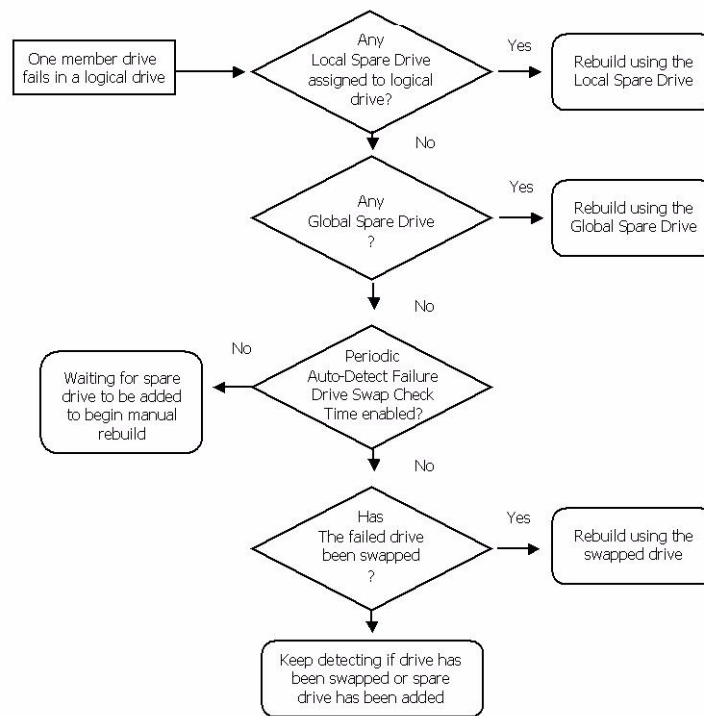


Fig. 1.13 Automatic Rebuild

Rebuild with Spare: When a member drive in a logical drive fails, the controller will first examine whether there is a Local Spare Drive assigned to this logical drive. If yes, rebuild is automatically started.

If there is no Local Spare available, the controller will search for a Global Spare. If there is a Global Spare, rebuild automatically begins using the Global Spare.

Failed Drive Swap Detected: If neither a Local Spare Drive nor Global Spare Drive is available, and the "Periodic Auto-Detect Failure Drive Swap Check Time" is "Disabled," the controller will not attempt to rebuild unless the user applies a forced-manual rebuild.

When the "Periodic Auto-Detect Failure Drive Swap Check Time" is "Enabled" (i.e., a check time interval has been selected), the controller will detect whether a faulty drive has been swapped (by checking the failed drive's channel/ID). Once the failed drive has been replaced with a healthy drive, the rebuild will begin immediately.

If the failed drive is not swapped but a local spare is added to the logical drive, rebuild will begin with the spare.

If the S.M.A.R.T. function is enabled on drives and the reaction scheme is selected for securing data on a failing drive, the spare will also be used for restoring data.

1.5.2

Manual Rebuild

When a user applies a forced-manual rebuild, the controller will first examine whether there is any Local Spare assigned to the logical drive. If yes, it will automatically start to rebuild. If there is no Local Spare available, the controller will search for a Global Spare. If there is a Global Spare, logical drive rebuild will be automatically conducted.

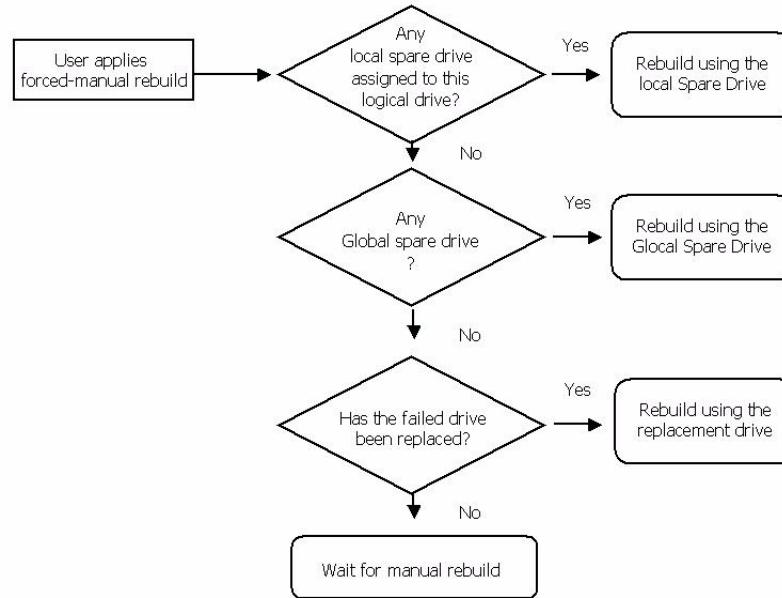


Fig. 1.14 Manual Rebuild

If none of the spares are available, the controller will examine the SCSI channel and ID of the failed drive. Once the failed drive has been replaced by a healthy one, it starts to rebuild using the new drive. If there is no available drive for rebuilding, the controller will not attempt to rebuild until the user applies another forced-manual rebuild.

1.5.3

Concurrent Rebuild in RAID (0+1)

RAID (0+1) withstands multiple drive failures with the condition that none of the mirrored pairs of data drives fails. (Two data drives mirroring each other must not fail together) Rebuilds can be automatically and concurrently conducted on more than one member if there is a sufficient number of pre-configured spare drives. Rebuild can also begin when the failed drives have been manually replaced. If failed drive detection has not been turned on, you may need to manually initiate the process.

1.6 Logical Volume (Multi-Level RAID)

1.6.1 What Is a Logical Volume?




Fig. 1.15 Logical Volume

A logical volume is a combination of RAID0 (Striping) and other RAID levels. Data written to a logical volume is first broken into smaller data segments and striped across different logical drives in a logical volume. Each logical drive then distributes data segments to its member drives according to its mirroring, parity, or striping scheme. A logical volume can be divided into a maximum of eight partitions. During normal operation, the host sees a non-partitioned logical volume or a partition of a partitioned logical volume as one single physical drive.

The benefits of using a logical volume are achieved by:

1. 1. Extending the MTBF (mean time between failure) by using more redundancy drives (spare drives).
2. 2. Decreasing the time to rebuild and reducing the chance of data loss caused by simultaneous drive failures because drives are included in different drive groups using a multi-level logical structure.
3. 3. Avoiding the chance of data loss caused by channel bus failure with flexible drive deployment.

As diagramed below, numerous drives can be included in a logical drive, and one of them is used for redundancy. By grouping these drives into several logical drives, and then into a logical volume, the chance of two drives failing in a logical unit is greatly reduced. Each logical drive can have one or more local spares. A failed drive can be immediately replaced by a local spare, reducing the risk of losing data if another drive fails soon afterwards.




Fig. 1.16 Logical Drive Composed of 24 Drives

As illustrated above, Configuration A is a RAID5 logical drive consisting of 24 physical drives. Configuration B is a logical volume made of four RAID5 logical drives.

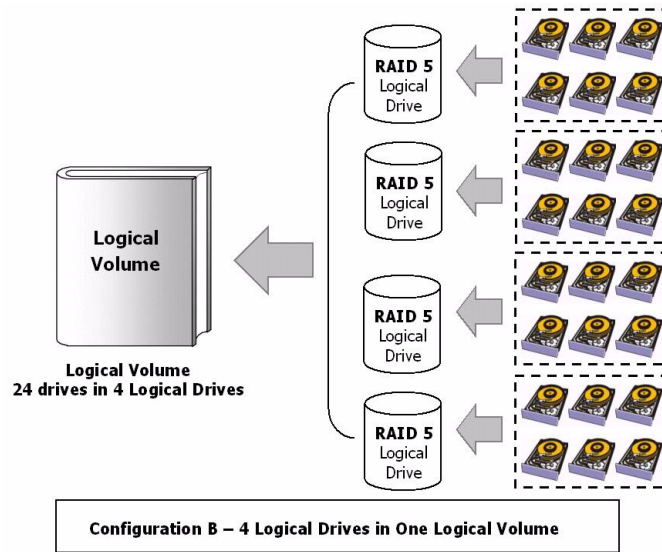


Fig. 1.17 Logical Volume Composed of 4 Logical Drives Formed by 24 Physical Drives

Configuration B can help reduce the chance of encountering points of failure:

- **Higher Redundancy:** Configuration A has one dedicated spare, while Configuration B allows the configuration of four spares. In Configuration B, the risk of simultaneous drive failure in a logical drive is significantly reduced compared to Configuration A. The total array capacity is comparatively smaller by the use of spares.
- **Less Rebuild Time:** The time during rebuild is a time of potential hazard. For example, a RAID5 logical drive can only withstand single drive failure - if another drive fails during the rebuild process, data will be lost. The time span for rebuilding a faulty drive should be minimized to reduce the possibility of having two drives fail at the same time. Configuration A is a large logical drive and takes a long time to rebuild. All members will be involved during the rebuild process. In Configuration B, the time span is shorter because only six members will participate when rebuilding any of the logical drives.
- **Channel Failure Protection:** Channel failure may sometimes result from unlikely matters such as a cable failure. A channel failure will cause multiple drives to fail at the same time and inevitably lead to a fatal failure. Using a logical volume with drives coming from different drive channels can avoid this point of failure.

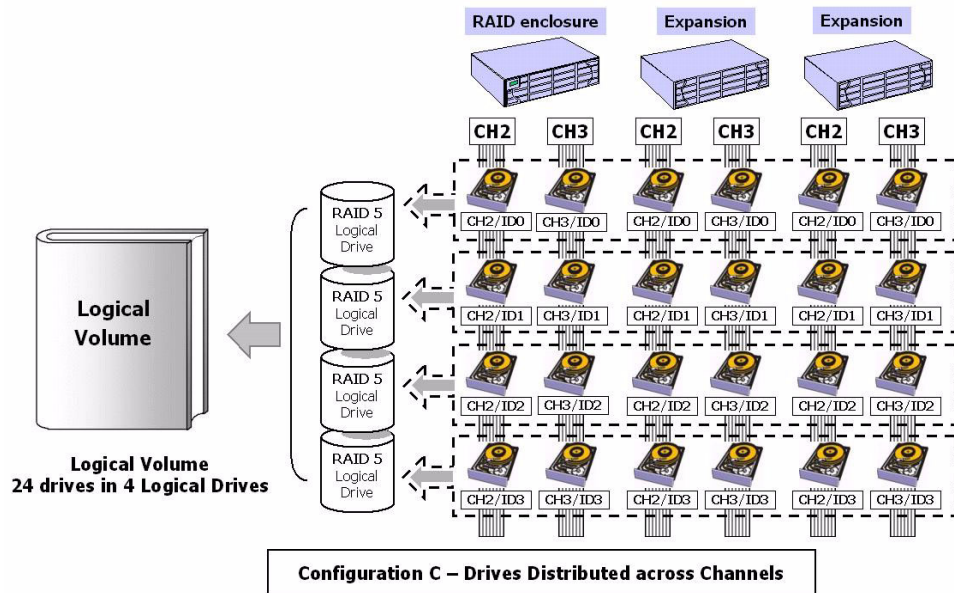


Fig. 1.18 Logical Volume with Drives on Different Channels

As illustrated above, if one of the drive channels fails, each logical drive loses one of its members. Logical drives still have the chance to rebuild its members. Data remains intact and the rebuild can be performed after the failed channel is recovered. No access interruptions to the logical volume will be experienced from the host side.

1.6.2

Can Spare Drives Assigned to a Logical Volume?

A Local Spare cannot be assigned to a Logical Volume. If a drive fails, it fails as a member of a logical drive; therefore, the controller allows a Local Spare's assignment to logical drives rather than logical volumes.

1.6.3

Limitations

The logical volume can not have any logical drive stated as "fatal failed." If there is any failed drive in any of its member logical drives, the controller will start to rebuild that logical drive. If any of the member logical drives fail fatally, the logical volume fails fatally and data will not be accessible.

To avoid a logical volume failure:

1. Logical drives as members of a logical volume should be configured in RAID levels that provide redundancy, i.e., - RAID levels 1 (0+1), 3 or 5.
2. Rebuild the logical drive as soon as possible whenever a drive failure occurs. Use of local spares is recommended.
3. A logical drive should be composed of physical drives from different drive channels. Compose the logical drive with drives from different drive channels to avoid the fatal loss of data caused by bus failure.

1.6.4

Partitioning the Logical Drive or Partitioning the Logical Volume?

Once a logical drive has been divided into partitions, the logical drive can no longer be used as a member of a logical volume. The members of a logical volume should have one partition only with the entire capacity.

If you want to use a partitioned logical drive for a logical volume, delete the other partitions in this logical drive until there remains one partition only with the entire capacity. Note that deleting the partition of the logical drive will also destroy all data. Data should be backed up before making partition configurations.

When a logical drive is used as a member of a logical volume, this logical drive can no longer be partitioned in “View and Edit Logical Drives.” Instead, the Logical Volume can be divided into eight partitions in “View and Edit Logical Volume.”

The procedure for partitioning a logical volume is the same as that for partitioning a logical drive. After the logical volume has been partitioned, map each partition to a host ID/LUN to make the partitions available as individual drives.

As members of a logical volume, all logical drives will be forced to adopt a consistent write policy. Whenever the write policy of a logical volume is changed, for example, the corresponding setting in its members will also be changed.

1.6.5

RAID Expansion with Logical Volume?

The Logical Volume can also be expanded using the RAID expansion function. The concept of expanding a logical volume is similar to that of expanding a logical drive. To perform RAID expansion on a logical drive, replace each member physical drive with a drive of larger capacity or add a new drive, then perform logical drive expansion to utilize the newly added capacity.

To perform RAID expansion on a logical volume, expand each member logical drive, then perform “RAID Expansion” on the logical volume.

Steps to expand a Logical Volume:

1. Expand each member logical drive.
2. Expand the logical volume.
3. Map the newly added capacity (in the form of a new partition) to a host LUN.



NOTICE!

- If a logical unit has already been partitioned, and you wish to expand its capacity, the added capacity will be appended to the last partition. You will not be able to proceed with expansion using a firmware version earlier than 3.27 when the unit already has eight partitions.
- Unless you move your data and merge two of the partitions, you will not be allowed to expand your logical volume. This is a precautionary limitation on logical unit expansion.

1.6.6

A Logical Volume with Logical Drives of Different Levels?

Multi-level RAID Systems

RAID (0+1) - this is a standard feature of Bosch RAID controllers. It brings the benefits of RAID1 (high availability) and RAID0 (enhanced I/O performance through striping). Simply choose multiple drives (the number must be even; i.e., a multiple of two) to compose a RAID1 logical drive, and RAID (0+1) will be automatically implemented.

RAID (1+0) or RAID 10 - a logical volume is a multi-level RAID implementation. A logical volume is a logical composition, which stripes data across several logical drives (the RAID0 method). A logical volume with several RAID1 members can be considered as a RAID (1+0), or RAID 10 volume.

RAID (3+0) or RAID 30 - a logical volume with several RAID3 members.

RAID (5+0) or RAID 50 - a logical volume with several RAID5 members.

2 RAID Planning

This chapter summarizes the RAID configuration procedures and provides some useful tools for first-time configuration.

2.1 Considerations

By the time you understand the basic ideas behind RAID levels, you may still be wondering how to begin. Here are the answers to some questions that may help you through the decision making process.

2.1.1 How many physical drives do you have?

When initially creating the drive groups, you should know how many drives you have in your RAID system and in the expansion enclosure(s) or JBOD(s) attached to the RAID controlling unit.

2.1.2 Limitations? Does my OS or HBA support 64-bit LBA?

A maximum of 64TB capacity can be included in a single logical drive.

If you have a picture of large arrays in mind, make sure both your host OS and HBA support 64-bit LBA. Not all OSes and HBA cards support 64-bit LBA. Some OSes, such as SuSE Linux 9, Windows 2003, Linux kernel 2.6X, support the addressability. Your HBA must also support this feature. Please refer to the documentation that came with your OS and HBA card to ensure the support for this feature.

Up to 128 members can be included in each logical drive. However, extreme array sizes can cause operational problems with system backup or management and should be avoided.

2.1.3 How many drives on each drive channel?

Optimal system planning is always a compromise between pros and cons. As a general rule, the number of drives you should connect on each channel equals the data bus bandwidth divided by the maximum transfer rate you can get from each of your hard drives. Knowing the mechanical performance of your hard drives can also help to determine how many drives to connect over a drive channel.

Always use fast and large drives of the same capacity to compose a logical array. A logical drive composed of an adequate number of larger drives can be more efficient than a logical drive comprised of many but smaller drives.

2.1.4 How many drives would you like to appear to the host computer?

You must decide what capacity will be included in a logical configuration of drives, be it a logical drive, a logical volume, or a logical partition. A logical configuration of drives will appear to the host as a single capacity volume.

You may compose a large logical volume consisting of drives on different drive channels/enclosures, and have it divided into smaller partitions. These partitions are then associated with host channel ID/LUNs. Each logical partition will appear as an independent capacity volume.

2.1.5 What kind of host application?

The frequency of read/write activities can vary from one host application to another. The application can be a SQL server, Oracle server, Informix, or other database server of a transaction-based nature. Applications like video playback and video post-production editing require read/write activities of large files coming in a sequential order.

Choose an appropriate RAID level for what is the most important for a given application – capacity, availability, or performance. Before creating your RAID, you need to choose an opti-

mization scheme and optimize each array/controller for your application. Stripe size and write policy can be adjusted on a per logical drive basis.

2.1.6

Multi-path, hub, or switch?

Unpredictable situations like a cable coming loose can cause system down time. Fibre Channel redundant path configurations using a flexible LUN mapping method can guarantee there is no single point of failure. You may avail a logical drive on two separate host channels and let the multi-path software running on the host manage the access to this logical drive. The use of a Fibre Channel switch makes cabling and topology more flexible. Change host-side protocol and other associated settings to adjust the RAID array to your demands.

2.1.7

Optimization mode




Fig. 2.1 Optimization Setting

You should select an optimization scheme best suited to your applications before configuring a RAID array. Once the optimization mode is selected, it will be applied to all arrays in the system.

Two options are available: Sequential I/Os and Random I/Os.

You may refer to *Section 5.2 Caching Parameters* and *Section 7.2 Caching Parameters* for the stripe size variables and their relation to RAID levels.

Numerous controller parameters are tuned for each optimization mode. Although stripe size can be adjusted on a per logical drive basis, users are not encouraged to make a change to the default values.

For example, smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. However, using the wrong stripe size can cause problems. When an array of the 4KB stripe size receives files of 128KB size, each drive will have to write many more times to store data fragments of the size of 4KB.

The default values in optimization modes guarantee the optimal performance for most applications. For all the controller parameters that are related to system performance and fault-tolerance see *Section 2.4 Tunable Parameters*.

2.1.8

What RAID level?

Different RAID levels provide varying performance features and fault tolerance levels. For explanation of RAID levels see *Section 1 RAID Functions: Introduction*.

Data availability comes with the tradeoff that more processing power is consumed because the RAID subsystem needs to generate/retrieve parity data. For example, in a performance-oriented configuration without availability concerns, you may configure an adequate number of drives into one or several RAID0 logical drives to maximize the array performance.

2.1.9

Any spare drives?

(Swap Drive Rebuild / Spare Drive Rebuild)

Spare drives allow for the unattended rebuilding of a failed drive, heightening the degree of fault tolerance. If there is no spare drive, data rebuild has to be manually initiated by replacing a failed drive with a healthy one.

As is often ignored, a spare drive (whether dedicated or global) must have a capacity no smaller than the members of a logical drive.

2.2 Configuring the Array

2.2.1 Starting a RAID System

Here is a flowchart illustrating basic steps to be taken when configuring a RAID system. Hardware installation should be completed before any firmware configuration takes place.




Fig. 2.2 Array Configuration Process

Drives must be configured and the subsystem properly initialized before a host computer can access the storage capacity.

1. Use the LCD keypad panel, RS-232C terminal program, or the RAIDWatch manager to start configuring your array.
2. When powered on, the controller scans all the hard drives that are connected through the drive channels. If a SCSI hard drive is connected after the controller completes initialization, use the "Scan SCSI Drive" function (only on SCSI-based controllers or subsystems) to let the controller recognize its presence.
3. Optimize the controller's parameters for your host applications.
4. Configure one or more logical drives to contain your hard drives based on the requirements of the desired RAID level, and/or partition the logical drive or logical volume into two or more partitions.

NOTICE!

- A "logical drive" is a set of drives grouped together to operate under a given RAID level and it appears as a single contiguous volume. The controller is capable of grouping drives into as many as 128 logical drives, configured in the same or different RAID levels.
- A maximum of 64 "logical volumes" can be created each from one or several logical drives. A logical drive or logical volume can be divided into a maximum of 64 "Partitions."

1. The next step is to make logical drives or storage partitions available through the host ports. When associated with a host ID or LUN number, each capacity volume appears as one system drive. The host SCSI or Fibre adapter will recognize the system drives after the host bus is re-initialized.
2. The last step is to save your configuration profile in the host system drive (the use of RAIDWatch manager is necessary for this purpose) or to the logical drives you created. The subsystem is independent from the host operating system. The host operating system will not be able to tell whether the attached storage is a physical hard drive or the virtual system drives created by the RAID subsystem.



2.3 Operation Theory

2.3.1 I/O Channel, Host ID, and LUN

Depending on the interface used by a RAID system, a SCSI channel (channel bus) can connect up to 15 drives (excluding the RAID controller itself). A Fibre Channel supports up to 126 drives in a loop. Depending on the subsystem design, a SCSI-based subsystem may have eight (8) disk drives connected through a drive channel while a SATA-based subsystem has one channel dedicated to each of its disk drive. Each device occupies one unique ID. The illustration belows shows the idea of mapping a system drive to host ID/LUN combinations (see *Figure 2.3*). The host ID is like a cabinet, and the drawers are the LUNs (LUN is short for Logical Unit Number). Each cabinet (host ID) contains up to 32 drawers (LUNs). Data can be made available through one of the LUNs of a host ID. Most host adapters treat a LUN like another device.




Fig. 2.3 SCSI ID/LUNs

2.3.2 Grouping Drives into an Array




Fig. 2.4 Connecting Drives

The logical view of the connection between RAID controller(s) and disk drives should be similar to the one shown above. Drives are connected through I/O paths on a back-end PCB that have been designated as drive channels.

The next diagram shows two logical configurations of drives and the physical locations of its members. There is no limitation on the locations of spare drives or drive members.

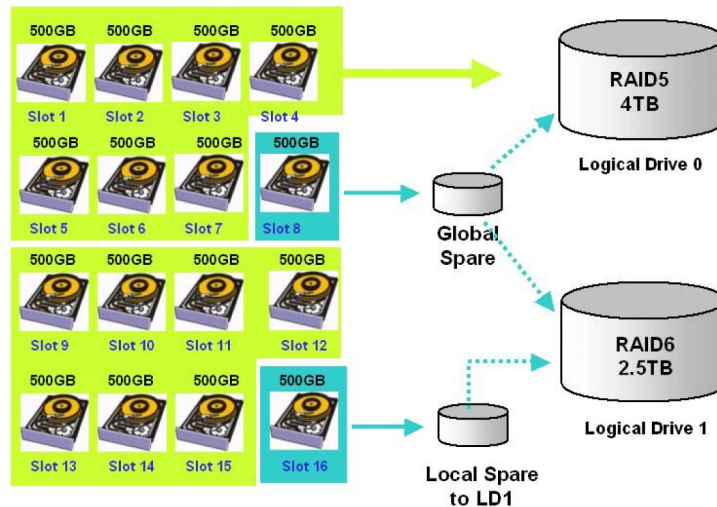


Fig. 2.5 Physical Locations of Drive Members

A drive can be assigned as the Local Spare Drive that serves one specific logical drive, or as a Global Spare Drive that participates in the rebuild of any logical drive. Spares automatically joins a logical drive in order to replace a failed member when a drive fails. Spares are not applicable to logical drives that have no data redundancy (NRAID and RAID 0).

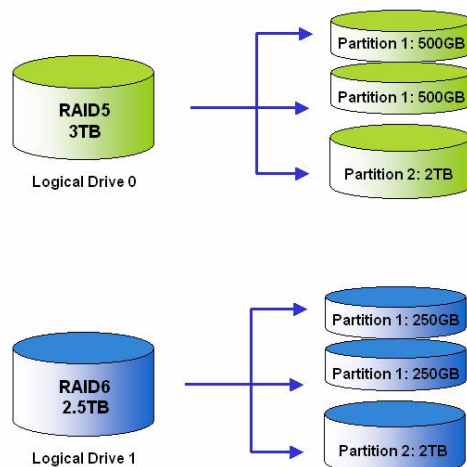


Fig. 2.6 Partitions in Logical Configurations

You may divide a logical drive or logical volume into partitions of desired capacity, or use the entire capacity as a single volume.

1. It is not a requirement to partition any logical configuration. Partitioning helps to manage a massive capacity.
2. Note that a logical drive cannot be included in a logical volume if it has already been partitioned.

2.3.3 Making Arrays Available to Hosts




Fig. 2.7 Mapping Partitions to Host ID/LUNs

Host ID mapping is a process that associates a logical configuration of drives with a host channel ID/LUN. To avail logical partitions on host channel(s), map each partition to a host ID or one of the LUNs under host IDs. Each ID or LUN will appear to the host adapter as one virtual hard drive.

There are alternatives in mapping for different purposes:

1. Mapping a logical configuration to IDs/LUNs on different host channels allows two host computers to access the same array. This method is applicable when the array is shared in a clustering configuration. Multi-path management software is required on the host side to avoid access contention.
2. Mapping partitions of an array to IDs/LUNs across separate host channels can distribute workload over multiple data paths.
3. Mapping across separate host channels also helps to make use of all bandwidth in a multi-path configuration. Firmware automatically manages the process when one data path fails and the workload on the failed data path has to be shifted to the existing data paths.

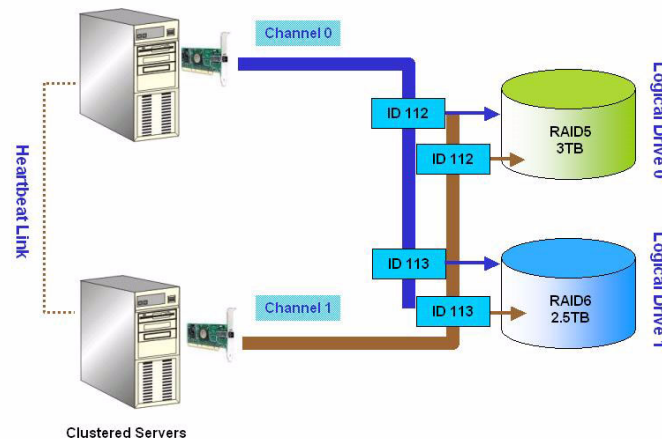


Fig. 2.8 Mapping in a Multi-path Environment

2.4 Tunable Parameters

Fine-tune the subsystem and the array parameters for your host applications. Although the factory defaults guarantee the optimized operation, you may refer to the table below to facilitate tuning of your array. Some of the performance and fault-tolerance settings may also be changed later during the preparation process of your disk array.

Use this table as a checklist and make sure you have each item set to an appropriate value.

(1)	Parameters that should be configured at the initial stage of system configuration
(2)	Parameters that can be changed later
(3)	Non-critical

Table 2.1 Controller Parameter Settings

User-Defined Parameters	Default	Alternate Settings
Fault Management:		
(1) Automatic Logical Drive Rebuild - Spare Drive	Enabled when Spare Drive is available	RAID 1 + Local Spare RAID 3 + Local Spare RAID 5 + Local Spare Global Spare
(1) S.M.A.R.T.	Detect Only	Disabled Perpetual Clone Clone + Replace
(3) Clone Failing Drive	Manual function	Replace After Clone Perpetual Clone
(1) Rebuild Priority	Normal	Low Improved High
(1) Verification on Write	Disabled	On LD Initialization On LD Rebuild On Normal Drive Writes
(3) SDRAM ECC	Enabled	Disabled
(2) Periodic Cache Flush	Disabled	Continuous to 10 minutes
(1) Event Notification	Reports to user interface and onboard alarm	Over SNMP Traps Over Java-Based Management Software
(2) Periodic Auto-Detect Failure Drive Swap Check Time	5	Disabled, 5 to 60 seconds
(2) Periodic Drive Check Time	30	Disabled, 0.5 to 30 seconds
(2) Rebuild Priority	normal	Low, normal, improved, high

Controller:		
(1) Channel Mode	*	Host, Drive, RCCOM, Drive + RCCOM
(1) Host and Drive Channel IDs	*	*
(1) Controller Unique Identifier	Preset on some models	Hex number from 0 to FFFFF (FW 3.25 and above)
(2) Data Rate	Auto	Depends on problems solving
(1) Date and Time	N/A	
(1) Time Zone	+ 8 hrs	

Optimization Mode:		
(1) Write-back Cache	Enabled	Disabled
(1) Array Stripe Size	Related to controller general setting	4KB to 1024KB
(2) Adaptive Write Policy	Disabled	Enabled
(1) Optimization for Random/Sequential	Sequential	Either
(2) Array Write Policy	Related to controller general setting	W/B or W/T

Host- and Drive-side Parameters:		
(1) Data Transfer Rate	*	Host Side: Asynchronous to 4GHz Drive Side: Asynchronous to 3GHz
(1) Maximum Tag Count	16	1-128
(1) Maximum Queued I/O Count	256	1 to 1024
(2) LUN's per ID	8	Up to 32
(1) Periodic Drive Check Time	30	Enabled
(1) Periodic SAF-TE and SES Device Check Time	30	Disabled to 60 seconds
(1) Periodic Auto-Detect Failure Drive Swap Check Time	5	5 to 60 seconds
(1) Number of Concurrent Host-LUN Connection	4	1 to 1024

Host- and Drive-side Parameters:		
(1) Tags per Host-LUN Connection	32	1 to 256
(1) Wide Transfer	*	Enabled/Disabled
(1) Drive I/O Timeout	30	0.5 to 30 seconds
Tag Command Queuing Native Command Queuing	16	1 to 128
(3) Drive Spindown Idle Delay Period	Disabled	Enabled
(3) Drive Delayed Write	Enabled	Disabled

Spin-Up Parameters: (available on SATA-based models)		
(1) Disk Access Delay Time	25 seconds	No Delay, 5 to 75 seconds

Data Integrity:		
(3) Task Scheduler	N/A	Execute on initialization Start time and date Execution period Media scan mode Media scan priority Logical drive selection

Array Configuration		
(1) Disk Reserved Space	256MB	
(1) AV Optimization Mode	Disable	Fewer Streaming Multiple Streaming
(1) Max Drive Response Timeout	Disabled	160, 320, or 960ms
(2) Array Assignment	Primary controller	Secondary controller
(1) Array Partitioning	1	Up to 64
(1) Auto-assign Global Spare	disabled	enabled

Enclosure Monitoring:		
(2) Periodic SAF-TE/ SES Device Check Time	30 seconds	Disabled, 50ms~60 seconds
(2) Event Triggered Operation	N/A	Controller, fan, PSU, BBU, and elevated temperature Auto-shutdown: 2 mins~1 hour
(1) Thresholds for Volt- age and Temperature Self-Monitoring	CPU temp: 0~90C Board temp: 0~80C 3.3V: 2.9~3.6V 5V: 4.5~5.5V 12V: 10.8~13.2V	User-defined; do not change parameters unless necessary

Others:		
(3) Password	N/A	User-Defined; Password Validation Timeout: 1 second to Always Check Configurable
(3) LCD Display Controller Name	N/A	User-defined
(1) Network Protocol Support	All enabled except SSH	No configuration options
(1) Cylinder/Head/ Sector Mapping	N/A	Depends on host OS

Table 2.2 User-Defined Parameters

3 Accessing the Array: Serial Port, Ethernet, and Access Security

3.1 RS-232C Serial Port

Bosch's controllers and subsystems can be configured via a PC running a VT-100 terminal emulation program, or a VT-100-compatible terminal. RAID enclosures usually provide one or more DB-9 or audio jack RS-232C ports. Simply use an RS-232C cable to connect between the subsystem/controller's RS-232C port and the PC serial (COM) port.

Make sure you use the included null modem (IFT-9011) to convert the serial port signals. A null modem is always provided with your RAID subsystem/controller. The null modem has the serial signals swapped for connecting to a standard PC serial interface.

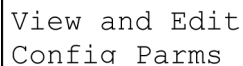
The following are guidelines on using the serial port:

- The serial port's default is set at 38400 baud, 8 bit, 1 stop bit and no parity. Use the COM1 serial port of the controller for terminal management session.
- In most cases, connecting RD, TD, and SG is enough to establish communication with a terminal.
- If you are using a PC as a terminal, any VT-100 terminal emulation software will suffice. Microsoft® Windows includes a terminal emulation program as presented with the "(Hyper) Terminal" icon in the Accessories window.
- For other details on connecting to the serial port, please refer to the Installation Manual that came with your subsystem.

3.1.1 Configuring RS-232C Connection via LCD Keypad Panel

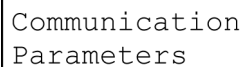
The subsystem's baud rate default is set to 38400. If necessary, follow the steps below to change the baud rate using the front panel keypad:

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Configuration ..", then press ENT.



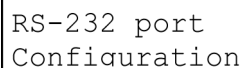
View and Edit
Config Params

Select "Communication Parameters ..", then press ENT.



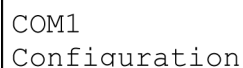
Communication
Parameters

Select "RS-232 Configuration ..", then press ENT.



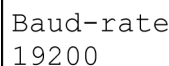
RS-232 port
Configuration

Select "COM1 Configuration ..", then press ENT.



COM1
Configuration

Select "Baud-rate XXXX..", then press ENT.



Baud-rate
19200

The baud rate default is 38400. If another baud rate is preferred, press the up or down arrow keys to select the baud rate, then press ENT for 2 seconds to confirm the selected baud rate. Set the identical baud rate to your RAID array and your terminal computer.

```
Baud-rate
38400
```

- The following baud rates are available: 2400, 4800, 9600, 19200, and 38400.
- Terminal connection should work properly using the above setting. You may check the following option in your COM port configuration if you encounter problems:
“Term Emul. Enab ..”: Make sure the terminal function has not been accidentally disabled.

3.1.2

Starting RS-232C Terminal Emulation

The keys used when operating via the terminal are as follows:

← → ↑ ↓	To select options
[Enter]	To go to a submenu or to execute a selected option
[Esc]	To cancel an option or return to the previous menu
[Ctrl]+[L]	To refresh the screen information




Fig. 3.1 RS232C Initial Screen



NOTICE!
If the RS-232C cable is connected while the controller is powered on, press [Ctrl] + [L] to refresh the screen information. Key combinations may vary depending on the terminal software you use.

The initial screen appears when the controller finishes its self-test and is properly initialized. Use ↑ ↓ arrow keys to select terminal emulation mode, then press [ENTER] to enter the Main Menu.

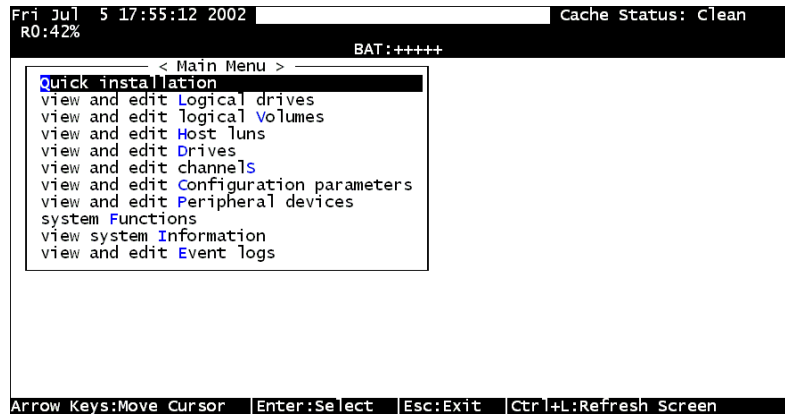


Fig. 3.2 Functional Item Page

Choose a functional item from the main menu to begin configuring your RAID.

3.2 Communication Parameters

The Communication Parameters is the first functional sub-menu you access once the physical RS-232C link is connected. In addition to the baud rate and terminal emulation options, the sub-menu contains several other options to prepare your management session using the Ethernet connection.

To access the sub-menu, use your arrow keys to select “View and Edit Configuration Parameters,” and press Enter on “Communication Parameters.” This provides menu accesses to “RS-232 Port Configuration” and “Internet Protocol <TCP/IP>.”

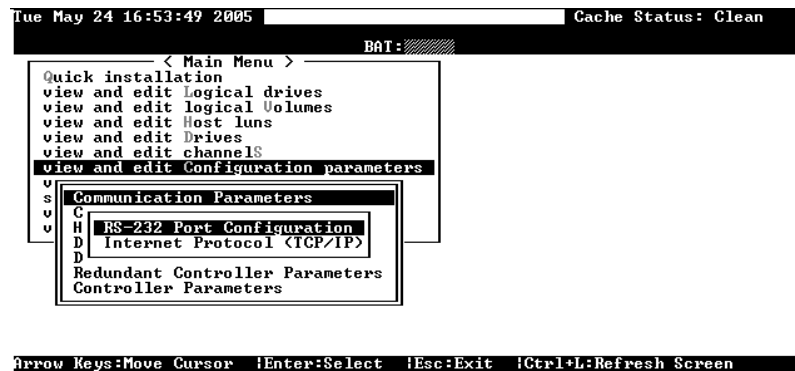


Fig. 3.3 Configuration Parameters Page

3.2.1

RS-232 Port Configuration

The “RS-232 Port Configuration” provides access to change the COM port operating parameters. Each COM port (COM1 or COM2) selection menu features two communication parameters: “Baud Rate” and “Terminal Emulation.”

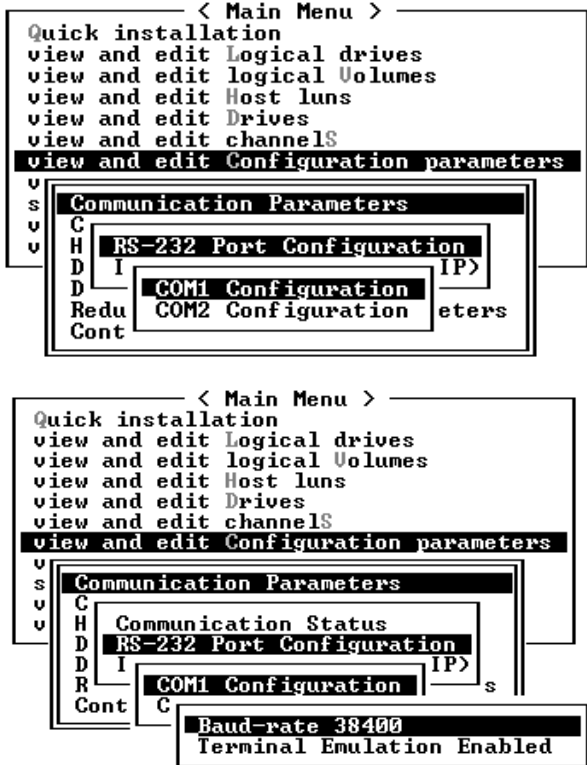


Fig. 3.4 RS-232 Port Configuration Page

3.2.2

Terminal Emulation:

The Terminal Emulation setting on the COM1 port is enabled by default. Usually there is no need to change this setting.

3.2.3

Baud Rate

To change the Baud Rate setting, select the highlighted menu item by pressing Enter. Available options will be displayed in a pull-down menu. Select by pressing Enter and press ESC several times to return to the previous configuration screen.




Fig. 3.5 Baud Rate Setting Page

3.2.4 Internet Protocol <TCP/IP>

The Internet Protocol menu allows you to prepare the management access through the subsystem/controller's RJ-45 Ethernet port.

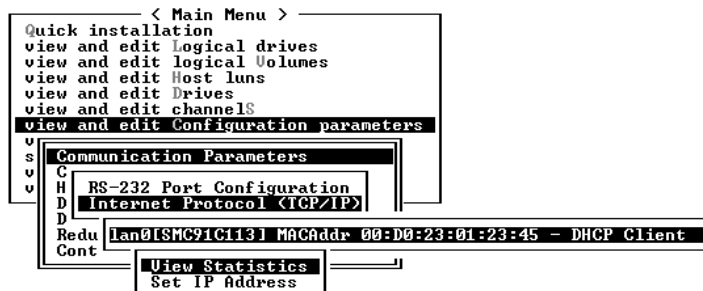


Fig. 3.6 Internet Protocol Page

To access the configuration options, press Enter on “Internet Protocol <TCP/IP>” to display the information of Ethernet port. Press [ENTER] on the chip information to display the “View Statistics” and the “Set IP Address” options.

3.2.5 View Statistics




Fig. 3.7 View Statistics Page

3.2.6 Set IP Address




Fig. 3.8 Internet Protocol Page

Provide a valid IP address for your subsystem/controller's Ethernet port. Consult your network administrator for an IP address and the associated NetMask and Gateway values. You may also key in “DHCP” if your local network supports automatic IP configuration.

NOTICE!

One drawback of using DHCP is that if cable disconnection or other unpredictable network faults occur, your Ethernet port may be assigned with a different IP. This may cause problems for the management sessions using the RAIDWatch manager. You may not be able to receive important event messages. RAIDWatch and its sub-module, the Configuration Client (the event notification utility), recognize a configured array through its Ethernet port IP.



3.3 Out-of-Band via Ethernet

The RAIDWatch manager software provides a graphical interface to RAID subsystems. RAID-Watch comes with an installer program that facilitates the installation of software modules for local/remote access to the array through standard TCP/IP.

Below is a tree structure of RAIDWatch's installation options:




Fig. 3.9 Installation Options

If you install RAIDWatch using the “*Typical*” option or the first two custom modes “*Centralized Management*” and “*Stand-alone (on Host)*,” prepare the following:

1. TCP/IP for the network connection to the controller/subsystem's Ethernet port.
2. Configure a valid IP and appropriate network settings for the RAID subsystem/controller's Ethernet port.

If you install RAIDWatch using the third installation option, “*Stand-alone (on Subsystem)*,” prepare the following:

1. TCP/IP for the network connection to the subsystem/controller's Ethernet port.
2. Configure a valid IP and appropriate network settings for the RAID subsystem/controller's Ethernet port.
3. Array reserved space: If you select the “Custom” installation mode, “Stand-alone (on Subsystem),” the RAIDWatch installation program will automatically create a reserved space on the first disk drive and distribute necessary program files to that space.
4. Unless the in-band connection over SCSI host busses is preferred, there is no need to configure the Peripheral Device setting with the later versions (2.1 and later) of RAID-Watch manager using Ethernet connection.

NOTICE!

- Only the connections via Ethernet are discussed here.
- Unless you want to start using RAIDWatch under the following conditions, you must not configure the Peripheral Device settings:
 - Using RAIDWatch to access a brand new array that has no configured drive groups.
 - Using the In-band method (over the existing host data paths) for array management.
 The Peripheral Device settings present the RAID array on the host data paths as if it is an external device.

If you use an Ethernet link and then unnecessarily configure the Peripheral Device settings, the application servers will mistakenly identify the RAID array as an external device and request a driver. Normal RAID array operation does not require supporting drivers.



3.3.1 What is the “Disk Reserved Space?”

If you install RAIDWatch using the third installation option, “Stand-alone (on Subsystem),” user’s configuration data and the manager’s main programs will be kept in a small section of disk space on a data drive.

After you create a logical drive, all its member drives will have a reserved space. The segregated disk space is called a “Disk Reserved Space.” When configuring a logical drive, the firmware automatically segregates 256MB of disk space from each of the member drives. The reserved space also allows “Drive Roaming.” Arrays (hard drives included in logical configurations) can be readily recognized by another RAID subsystem/RAID controller. This is useful when hardware faults occur, and you have to replace a faulty subsystem/controller.

3.3.2 Other Concerns

Availability Concern

The reserved space information can be seen during the array creation process or in the “View and Edit Drives” menu.

Whatever data is put into the reserved space, the firmware will automatically duplicate and distribute it to the reserved section on every data drive. Even if one hard drive or one logical drive fails, an exact replica still resides on other member drives.

Web-based Management

When firmware is installed on the stand-alone controller, once properly configured, the subsystem/controller’s Ethernet port behaves like an HTTP server. This applies to RAIDWatch’s third installation mode, “Stand-alone (on Subsystem).”

3.3.3 Connecting Ethernet Port:

Use a LAN cable to connect the Ethernet port(s) on the subsystem’s RAID controller unit(s). Connect the cables between subsystem/controller’s Ethernet port and an Ethernet port from your local network.

3.3.4 Configuring the Controller

To prepare the subsystem/controller for Ethernet connection:

1. Use a Terminal Emulator to Begin Configuration
Connect the subsystem’s serial port to a PC running a VT-100 terminal emulation program or a VT-100-compatible terminal.
Make sure the included null modem is already attached to enclosure serial port or the management computer’s COM port. The null modem converts the serial signals for connecting to a standard PC serial interface. For more details, please refer to the descriptions above in *Section 3.1 RS-232C Serial Port* and *Section 3.2 Communication Parameters*.
2. Assign an IP Address to the Ethernet Port
Assign an IP address to the controller Ethernet port and specify the Net Mask and Gateway values. Power off your system and then power on again for the configuration to take effect.
Select "View and Edit Configuration Parameters" from the main menu. Select "Communication Parameters" -> "Internet Protocol (TCP/IP)" -> press [ENTER] on the chip hardware address -> and then select "Set IP Address."
You may also use an auto discovery protocol such as DHCP. Simply key in “DHCP” in the IP address field.

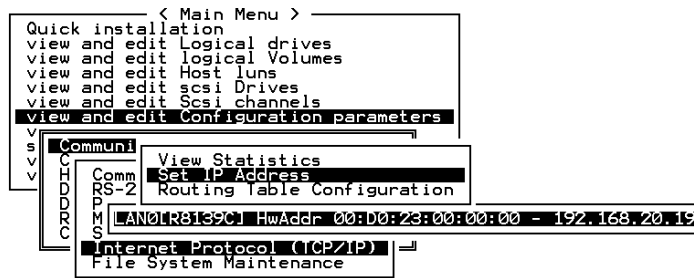


Fig. 3.10 Set IP Address Page

Provide the IP address, NetMask, and Gateway values accordingly.




Fig. 3.11 Provide IP Address Page

PING the IP address from your management computer to make sure the link is up and running.

3.4

Telnet Connection

1. Use an Ethernet cable with RJ-45 phone jacks to connect the Ethernet port on the sub-system/controller module.
2. Connect the other end of the Ethernet cable to your local area network. An IP address should be acquired for the subsystem's Ethernet port. The subsystem firmware also supports automatic client configuration such as DHCP.
3. Consult your network administrator for an IP address that will be assigned to the sub-system/controller Ethernet port.
4. Assign an IP address to the subsystem/controller's Ethernet port and specify the Net-Mask and Gateway values. Power off your system and then power on again for the configuration to take effect.
5. Select "View and Edit Configuration Parameters" from the main menu on the terminal screen. Select "Communication Parameters" -> "Internet Protocol (TCP/IP)" -> press ENTER on the chip hardware address -> and then select "Set IP Address."
6. Provide the IP address, NetMask, and Gateway values accordingly.
7. PING the IP address from your management computer to make sure the link is up and running.
8. Open a command prompt and key in "telnet <IP address>" to access the embedded firm-ware utility.

3.5 Secure Link over SSH

Firmware supports remote management over the network connection and the security under SSH (Secure Shell) protection. SSH is widely used for its ability to provide strong authentication and secure communications over insecure channels

To access the firmware embedded configuration utility over the SSH network, do the following:

1. Open a web browser.
2. Key in the line below in the browser's URL field.
`http://<controller IP>/ssh.html`
3. Enter controller name (if there is any) as user name and the password set for the RAID subsystem. If neither the controller name or password has been set, press Enter at each command prompt to proceed.

4 LCD Screen Messages

4.1 The Initial Screen




Fig. 4.1 LCD Initial Screen

Status/Data Transfer Indicator:

Ready	There is at least one logical drive or logical volume mapped to a host ID/LUN combination.
No Host LUN	No logical drive created or the logical drive has not yet been mapped to any host ID/LUN.
<div> </div>	Indicates the statistic of internal processing resources, not the host bus throughput. Each block indicates 256Kbytes of data being processed.

4.2 Quick Installation Screen

Bild einfügen: Quick_Installation_Screen



Fig. 4.2 Quick Installation Screen

Press [ENT] to create a logical drive. The subsystem/controller will start initialization of one logical drive to include all connected disk drives and automatically map the logical drive to the first ID and LUN number of the first host channel. The “Quick Installation” can only be performed when there is no logical drive.

4.3 Logical Drive Status




Fig. 4.3 Logical Drive Status Screen

Logical Drive:	The Logical Drive number.
RAID level:	The RAID level applied for this logical drive.
Number of Drives:	The number of physical drives included in this configuration.

Logical Drive status:	
XxxxMB	The capacity of this logical drive.
SB=x	Standby drives available to this logical drive. Except the spares dedicated to other logical configurations, all spare drive(s) will be counted in this field, including Global and Local Spares.
xxxxMB INITING	The logical drive is now initializing.
xxxxMB INVALID	For firmware version before 3.31: The logical drive has been created with "Optimization for Sequential I/O", but the current setting is "Optimization for Random I/O." -OR- The logical drive has been created with "Optimization for Random I/O," but the current setting is "Optimization for Sequential I/O." Firmware versions 3.31 and later have separate settings for array optimization and array stripe size. This message will not appear when the optimization mode is changed.
xxxxMB GD SB=x	The logical drive is in good condition.
xxxxMB FL SB=x	One member drive failed in this logical drive.
xxxxMB RB SB=x	Logical Drive is rebuilding.
xxxxMB DRVMISS	One of the drives is missing.
INCOMPLETE ARRAY	Two or more drives failed in this logical drive.

4.4 Logical Volume Status




Fig. 4.4 Logical Volume Status Screen

Logical Volume:	The Logical Volume number.
DRV=x:	The number of logical drive(s) contained in this logical volume.
Logical Volume ID Status:	The unique ID number of the logical volume (randomly generated by firmware).
Logical Volume Status:	
xxxMB	The capacity of this logical volume.

4.5 Drive Status

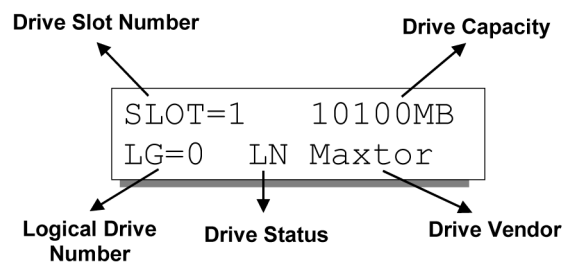


Fig. 4.5 Drive Status Screen

Drive Status:	
LG=x IN	Initializing
LG=x LN	On-line (already a member of a logical configuration)
LG=x RB	Rebuilding
LG=x SB	Local Spare Drive
GlobalSB	Global Spare Drive
NEW DRV	New drive
BAD DRV	Failed drive
ABSENT	Drive does not exist; an empty drive tray
MISSING	Drive missing (drive was once there)
SB-MISS	Spare drive missing

4.6

Channel Status

Host Channel




Fig. 4.6 Host Channel Status Screen

Drive Channel (Single Controller Configuration)




Fig. 4.7 Drive Channel Status Screen

Channel Mode:	
Host	Host channel mode
Drive	Drive channel mode
Default SCSI Bus Sync Clock:	
160M	The default setting of this channel is set to the 160Mhz mode
Async, 2.5 to 160Mhz	Data rate can be manually configured if the need arises
Available Host IDs:	
*	Multiple IDs applied (Host channel mode only)
(ID number)	Controller is using this ID for host LUN mapping.
NA	No ID applied

4.7 Controller Voltage and Temperature

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Peripheral Dev," then press ENT.

```
View and Edit
Peripheral Dev ↑
```

Press the up or down arrow keys to select "Ctrlr Peripheral Device Config..", press ENT and then choose "View Ctlr Periph Device Status..", then press ENT.

```
Ctrlr Peripheral
Device Config..
```

```
View Ctlr Periph
Device Status..
```

Press the up or down arrow keys to choose either "Voltage Monitor", or "Temperature Monitor".

```
Voltage Monitor
..
```

Select "Temperature and Voltage Monitor" by pressing ENT. Press the up or down arrow keys to browse through the various voltage and temperature statuses.

```
Temperature
Monitor ..
```

```
[+12V] 12.077V
Operation Normal
```

```
[+5v] 4.938v
Operation Normal
```

```
[CPU] 43.5°C
in Safe Range
```

```
[+3.3V] 3.384V
Operation Normal
```

```
[Board1] 46.5°C
in Safe Range
```

```
[CPU] 43.5°C
in Safe Range
```

```
[Board2] 46.5°C
in Safe Range
```

```
[+12v] 12.077v
Operation Normal
```

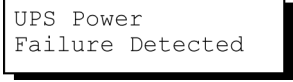
4.8 View and Edit Event Logs

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Event Logs," then press ENT.



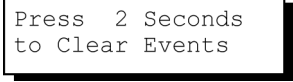
View and Edit
Event Logs ↑

Press the up or down arrow keys to browse through the existing event log items.



UPS Power
Failure Detected

To delete a specified item and all events prior to this event, press ENT for 2 seconds.



Press 2 Seconds
to Clear Events



NOTICE!

The event log will be cleared after the controller is powered off or reset.

5 LCD Keypad Operation

A navigation roadmap for the configuration menu options of LCD keypad is separately available as a PDF file. You may check your Product Utility CD or visit Bosch's Extranet site for the latest update.

5.1 Power On the RAID Enclosure

Before you start to configure a RAID system, make sure that hardware installation is completed before any configuration takes place. Power on your RAID enclosure.

5.2 Caching Parameters

5.2.1 Optimization Modes

Mass storage applications can be divided into two major categories according to its read/write characteristics: database and video/imaging. To optimize a subsystem for these two categories, two embedded optimization modes are available with system behaviors adjusted to different read/write parameters. They are the Optimization for Random I/O and the Optimization for Sequential I/O.

Limitations:

There are limitations on the use of optimization modes.

1. System default is Optimization for Sequential I/O.
2. You can select the stripe size of each array (logical drive) during the initial configuration. However, changing stripe size is only recommended for experienced engineers who have tested the effects of tuning stripe sizes for different applications.
3. The array stripe size can only be changed during the initial configuration process.
4. Once the controller optimization mode is applied, access to different logical drives in a RAID system will follow the same optimized pattern.

Database and Transaction-based Applications:

These kinds of applications usually include SQL server, Oracle server, Informix, or other database services. These applications keep the size of each transaction down to the minimum, so that I/Os can be rapidly processed. Due to their transaction-based nature, these applications do not read or write a bunch of data in sequential order - access to data occurs randomly. The transaction size usually ranges from 2K to 4K. Transaction performance is measured in "I/Os per second" or "IOPS."

Video Recording/Playback and Imaging Applications:

These kinds of applications usually includes video playback, video post-production editing, or other similar applications. These applications have the tendency to read or write large files from and into storage in sequential order. The size of each I/O can be 128K, 256K, 512K, or up to 1MB. The efficiency of these applications is measured in "MB/sec."

When an array works with applications such as video or image oriented applications, the application reads/writes from the drive as large-block, sequential threads instead of small and randomly accessed files.

The controller optimization modes have read-ahead buffer and other R/W characteristics tuned to obtain the best performance for these two major application categories.

5.2.2 Optimization Mode and Stripe Size

Each controller optimization mode has preset values for the stripe size of arrays created in different RAID levels. If you want a different setting for a subsystem with configured arrays, you have to backup or move the stored data and re-create the arrays:

1. Stop host I/O access.
2. Move or backup all of your stored data in the subsystem.
3. Change the optimization mode.
4. Reset the controller.
5. Re-create the array(s). Once the array(s) are created, stripe size cannot be changed.

Listed below are the default stripe sizes implemented with different optimization modes and RAID levels. These values should be sufficient for most applications:

	Opt. For Sequential I/O	Opt. for Random I/O
RAID0	128KB	32KB
RAID1	128KB	32KB
RAID3	16KB	4KB
RAID5	128KB	32KB
NRAID	128KB	32KB

Table 5.1 RAID Level, Optimization Modes, and Stripe Sizes

5.2.3 Optimization for Random or Sequential I/Os

```
Caching
Parameters  ..
```

```
Optimization I/O
Random
```

```
Optimization for
Sequential I/O?
```



NOTICE!

The original 512GB threshold on array optimization mode has been canceled. If the size of an array is larger than 16TB, only the optimization for sequential I/O can be applied. Logical drives of this size may not be practical; therefore, there is actually no limitation on the optimization mode and array capacity.

5.2.4 Write-back/Write-through Cache Enable/Disable

Write-back/Write-through Cache Enable/Disable

As one of the submenus in "Caching Parameters," this option controls the cached write function. Press ENT to enable or disable "Write-back Cache." Press ENT for two seconds to confirm. The current status will be displayed on the LCD.

```
Write-Back Cache
Enabled ..
```

The Write-through mode is safer if your subsystem is not managed by a controller pair and there is no battery backup.

```
Disable Write  
-Back Cache  ?
```

Write-back caching can dramatically improve write performance by caching the unfinished writes in memory and letting them be committed to drives latter in a more efficient manner. In the event of power failure, a battery module can hold cached data for days. In the event of controller failure in a dual-active configuration, data cached in the failed controller has an exact replica on its counterpart controller and therefore remains intact.

5.2.5

Periodic Cache Flush

If for a reason Write-Back caching is preferred for better performance and yet data integrity is of the concern, e.g., lack of the battery protection, the system can be configured to flush the cached writes at every preset interval.

Note that the “Continuous Sync” option holds data in cache for as long as necessary to complete a write operation and immediately commits it to hard drives if it does not come in a series of sequential writes.

Select from Main Menu “View and Edit Config Params,” “Caching Parameters,” and press ENT. Use the arrow keys to scroll through the options and select “Periodic CachFlush Time”, and then press ENT to proceed. The “Set Cache Flush Time – Disable” appears. The default is “Disable.” Use your arrow keys to select an option from “ConSync,” “1/2min,” to “10min.” “ConSync” stands for “continuously synchronized.” Press ENT to select and press ESC to leave and the setting will take effect immediately.

```
Period CachFlush  
Time - Disable
```

```
Set Cache Flush  
Time - Disable
```



NOTICE!

Every time you change the caching parameters, you must reset the controller for the changes to take effect.

5.3 View Connected Drives:

A RAID system consists of many physical drives that can be modified and configured as the members of one or several logical drives.

Press the front panel ENT button for two seconds to enter the Main Menu. Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Drives," then press ENT.

```
View and Edit  
Drives
```

Use the up or down arrow keys to scroll down the list of connected drives' information screens.

```
Slot=* 1010MB  
New DRV SEAGATE
```

You may first examine whether there is any drive installed but not shown here. If there is a drive installed but not listed, the drive may be defective or not installed correctly. Please check your enclosure installation and contact your system vendor.

Press ENT on a drive. Choose "View Drive Information" by pressing ENT. Use the up or down arrow keys to navigate through the screens.

```
View Drive  
Information ..
```

The Revision Number of the selected drive will be shown. Press the down arrow key to see other information.

```
Revision Number:  
0274
```

Other information screens include "Serial Number" and "Disk Capacity" (displayed in blocks; each block equals 512K Bytes).

NOTICE!

- Drives of the same brand/model/capacity might not feature the same block number.
- The basic read/write unit of a hard drive is a block. If members of a logical drive have different block numbers (capacity), the smallest block number will be taken as the maximum capacity to be used in every drive. Therefore, use drives of the same capacity.
- You may assign a Local/Global Spare Drive to a logical drive whose members have a block number equal to or smaller than the Local/Global Spare Drive but you should not do the reverse.



5.4 Creating a Logical Drive

To create a logical drive, press ENT for two seconds to enter the Main Menu. Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Logical Drives," and then press ENT.

```
View and Edit  ↓  
Logical Drives..
```

Press the up or down arrow keys to select a logical drive entry, then press ENT for two seconds to proceed. "LG" is short for Logical Drive.

```
LG=0  
Not Defined   ?
```

5.4.1 Choosing a RAID Level

Press the up or down arrow keys to choose the desired RAID level, then press ENT for two seconds. "TDRV" (Total Drives) refers to the number of all available drives.

```
TDRV=24 Create  
LG Level=RAID6 ?
```

5.4.2 Choosing Member Drives

Press ENT for two seconds; the message, "RAID X selected To Select drives", will prompt. Confirm your selection by pressing ENT.

```
RAID 6 Selected  
To Select drives
```

Press ENT, then use the up or down arrow keys to browse through the available drives. Press ENT again to select/deselect the drives. An asterisk (*) mark will appear on the selected drive(s). To deselect a drive, press ENT again on the selected drive. The (*) mark will disappear.

```
Slot=* 239367MB  
New DRV SEAGATE
```

After all the desired drives have been selected, press ENT for two seconds to continue.

```
SLOT=1 239367MB  
*LG=0 SL SEAGATE
```

5.4.3 Logical Drive Preferences

You may also choose "Change Logical Drive Parameter," then press ENT to change related parameters before initializing the logical drive.

```
Change Logical  
Drive Parameter?
```


5.4.4 Maximum Drive Capacity

Choose “Maximum Drive Capacity,” then press ENT. The maximum drive capacity refers to the maximum capacity that will be used in each member drive.

```
Maximum Drive  
Capacity      ..
```

If necessary, use the up and down arrow keys to change the maximum size that will be used on each drive.

```
MaxSiz= 239367MB  
Set to 239367MB?
```

5.4.5 Spare Drive Assignments

The Local (Dedicated) Spare Drive can also be assigned here. Press the up or down arrow keys to choose “Spare Drive Assignments,” then press ENT.

```
Spare Drive  
Assignments  ..
```

Available drives will be listed. Use the up or down arrow keys to browse through the drive list, then press ENT to select the drive you wish to use as the Local Spare Drive. Press ENT again for two seconds.

```
Slot=* 239367MB  
New DRV SEAGATE
```

```
Slot=* 239367MB  
*LG=0 SL SEAGATE
```

5.4.6 Disk Reserved Space

This menu allows you to see the size of disk reserved space. Default is 256MB.

```
Disk Rev. Space  
256MB      ..
```


5.4.7 Write Policy

This menu allows you to set the caching mode policy for this specific logical drive. “Default” is a neutral value that is coordinated with the subsystem’s caching mode setting. Other choices are “Write-back” and “Write-through.”

```
Write Policy  
Default     ..
```

5.4.8 Initialization Mode


This menu allows you to determine if the logical drive is immediately accessible. If the Online method is used, data can be written onto it before the array's initialization is completed. Users may proceed with array configuration, e.g., including this array in a logical volume.



Array initialization can take a long time especially for those comprising large capacity. Setting to "Online" means the array is immediately accessible and that the controller will complete the initialization when I/O demands become less intensive.

5.4.9 Stripe Size

This menu allows you to change the array stripe size. Setting to an incongruous value can severely drag performance. This item should only be changed when you can be sure of the performance gains it might bring you.



Listed below are the default values for an array. For each logical drive, the default value for the stripe size is determined by the Optimization Mode and the RAID level chosen.


	Opt. for Sequential I/O	Opt. for Random I/O
RAID0	128KB	32KB
RAID1	128KB	32KB
RAID3	16KB	4KB
RAID5	128KB	32KB
NRAID	128KB	32KB

Table 5.2 RAID Level, Optimization Modes, and Stripe Size

When you are done setting logical drive preferences, press ESC and use your arrow keys to select "Create Logical Drive?". Press ENT for two seconds to proceed.


5.4.10 Beginning Initialization

Press ESC to return to the previous menu. Use the up or down arrow keys to choose "Create Logical Drive," then press ENT for two seconds to start initializing the logical drive.



The Online Mode:

If the online initialization method is applied, the array will be available for use immediately. The array initialization runs in the background while data can be written onto it and users can continue configuring the RAID system.



The Offline Mode:

The controller will start to initialize the array parity if using the "offline" mode. Note that if NRAID or RAID 0 is selected, initialization time is short and completes almost immediately.

```
Initializing090%  
Please Wait!
```

```
LG=0 Initializat  
Ion Completed
```

The logical drive's information displays when the initialization process is completed. If the "online" mode is adopted, array information will be displayed immediately.

```
LG=0 RAID6 DRV=4  
478734MB GD SB=0
```

**NOTICE!**

Due to the operation complexity, the RAID Migration option is not available using the LCD keypad panel.

5.5 Creating a Logical Volume

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Volume," then press ENT.

```
View and Edit  
Logical Volume ↑
```

Press the up or down arrow keys to select an undefined entry for logical volume, then press ENT for two seconds to proceed. "LV" is short for Logical Volume.

```
LV=0  
Not Defined ?
```

Proceed to select one or more logical drives. Press ENT to proceed. "LD" is short for Logical Drive.

```
LV=0 Selected To  
Select LD Drives?
```

Use the up or down arrow keys to browse through the logical drives. Press ENT again to select/deselect the drives. An asterisk (*) mark will appear when the logical drive is selected. After all the desired logical drive(s) have been selected, press ENT for two seconds to continue.

```
LG0 RAID5 DRV=3  
2021MB GD SB=0
```

Two submenus will appear.

5.5.1 Initialization Mode

Array initialization can take a long time especially for those comprised of large capacity. Setting to "Online" means the array is immediately accessible and that the controller will complete the initialization when I/O demands become less intensive.

```
Initialization  
Mode
```

5.5.2 Write Policy

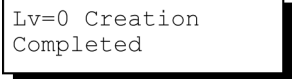
This menu allows you to set the caching mode policy for this specific logical volume. "Default" is a neutral value that is coordinated with the controller's caching mode setting. Other choices are "Write-back" and "Write-through."

```
Write Policy  
Default ..
```

When you are finished setting the preferences, press ENT for two seconds to display the confirm box. Press ENT for two seconds to start initializing the logical volume.

```
Create  
Logical Volume ?
```

The logical volume has been successfully created.



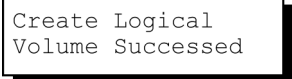
Lv=0 Creation
Completed

Press ESC to clear the message. Another message will prompt; press ESC to clear it.



Lv=0 ID=07548332
0024488MB DRV=2

Logical volume information will be displayed below.



Create Logical
Volume Succeeded

5.6 Partitioning a Logical Drive/Logical Volume

Partitioning and the creation of a logical volume, are not requirements for creating a RAID system. The configuration processes for partitioning a logical drive are the same as those for partitioning a logical volume.

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Volume," then press ENT.

```
View and Edit
Logical Volume ↓
```

Use the up or down arrow keys to select a logical volume, then press ENT.

```
LV=0 ID=685AE502
2021MB DRV=1
```

Use the up or down arrow keys to select "Partition Logical Volume," then press ENT.

```
Partition
Logical Volume..
```

The total capacity of the logical volume will be displayed as one partition. Press ENT for two seconds to change the size of the first partition.

```
LV=0   Prt=1:
      2021MB   ?
```

Use the up or down arrow keys to change the number of the flashing digit, (see the arrow mark) then press ENT to move to the next digit. After changing all the digits, press ENT for two seconds to confirm the capacity of this partition. You may also use arrow keys to move down to the next partition to configure more partitions.

```
LV=0   Part=0:
      2021MB
```

The rest of the drive space will be automatically allocated as the last partition. You may go on to create up to 16 partitions using the same method described above.

```
LV=0 Part=0:
      700MB   ?
```

```
LV=0 Partition=1
      1321MB   ?
```

Press ESC several times to go back to the Main Menu.



NOTICE!

If operated with a Unix-based system, reset the subsystem for the configuration to take effect if any changes have been made to partition sizes and partition arrangement.

5.7 Mapping a Logical Volume/Logical Drive to Host LUN

The process of mapping a logical drive is identical to that of mapping a logical volume. The process of mapping a logical volume is used as an example.

Press ENT for two seconds to enter the Main Menu. Use the up or down arrow keys to select

```
View and Edit  ↓
Host Luns     ..
```

The first available ID on the first host channel appears (usually channel 0).

```
CH=0 ID=0
..
```

Press the up or down arrow keys to select a configured host ID, and then press ENT for two seconds to confirm.

Press the up or down arrow keys to select the type of logical configuration to be associated with a host ID/LUN. Available choices are “Map to Logical Drive” or “Map to Logical Volume”. Confirm your choice by pressing ENT for two seconds.

```
Map to
Logical Drive ?
```

```
Map to
Logical Volume ?
```

Press the up or down arrow keys to select a LUN number, then press ENT to proceed.

```
CH0 ID0 LUN0 Not
Mapped
```

Press ENT for two seconds to confirm the selected LUN mapping.

```
Map Host LUN  ?
```

```
LV=0 ID=685AE502
478734MB DRV=1
```

Press the up or down arrow keys to select a logical drive, a logical volume, or a partition of them. Press ENT for two seconds to map the selected partition to this LUN. If the logical configuration has not been partitioned, you can map the whole capacity to a host LUN.

```
LV=0 PART=0
700MB ?
```

Press ENT for two seconds when prompted by “Map Host LUN” to proceed.

```
Map Host LUN
```

Mapping information will be displayed on the subsequent screen. Press ENT for two seconds to confirm the LUN mapping.

```
CH0 ID0 LUN0
MAP to LV0 PRT0?
```

The mapping information will appear for the second time. Press ENT or ESC to confirm, and the host ID/LUN screen will appear. Use the arrow keys to select another ID or LUN number to continue mapping your logical configurations or press ESC for several times to leave the configuration menu.

```
CH0 ID0 LUN0
MAPto LV0 PRT0
```

```
CH=0 ID=0
```

```
. .
```

With any of the host ID/LUN successfully associated with a logical capacity, the “No Host LUN” message in the initial screen will change to “Ready.”

If you want to create more host IDs, please see *Section 5.12 Viewing and Editing Channels*.

5.8 Assigning Spare Drive and Rebuild Settings

5.8.1 Adding a Local Spare Drive

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

```
View and Edit  
Drives ↓
```

Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a drive that is stated as "NEW DRV" or "USED DRV" that has not been included in any logical drive, nor specified as a spare drive or failed drive, then press ENT to select it.

```
Slot=* 1010MB  
NEW DRV SEAGATE
```

Press the up or down arrow keys to select "Add Local Spare Drive," then press ENT.

```
Add Local Spare  
Drive ..
```

Press the up or down arrow keys to select the logical drive where the Local Spare Drive will be assigned, then press ENT for two seconds to confirm.

```
LG0 RAID5 DRV=3  
2012MB GD SB=0
```

The message "Add Local Spare Drive Successful" will be displayed on the LCD.

```
Add Local Spare  
Drive Successful
```

5.8.2 Adding a Global Spare Drive

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

```
View and Edit  
Drives ↓
```

Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a disk drive that has not been assigned to any logical drive, then press ENT.

```
Slot=* 1010MB  
NEW DRV SEAGATE
```

Press the up or down arrow keys to select "Add Global Spare Drive," then press ENT.

```
Add Global Spare  
Drive ..
```

Press ENT again for two seconds to add the spare drive. The message, "Add Global Spare Drive Successful," will be displayed on the screen.

```
Add Global Spare  
Drive Successful
```

5.8.3

Rebuild Settings

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Config Params," then press ENT.

```
View and Edit
Config Params  ↑
```

Press the up or down arrow keys to select "Disk Array Parameters," then press ENT.

```
Disk  Array
Parameters..
```

Press the up or down arrow keys to select "Rebuild Priority Low," then press ENT. "Low" refers to the temporary default setting.

```
Rebuild Priority
Low              ..
```

Press ENT again and the abbreviation mark ".." will change to a question mark "?". Press the up or down arrow keys to select priority "Low," "Normal," "Improved," or "High".

```
Rebuild Priority
Low              ?
```

Press ENT to confirm and the question mark "?" will turn into "..".

```
Rebuild Priority
High            ..
```



NOTICE!

The rebuild priority determines how much of the system resources are used when rebuilding a logical drive. The default setting of the rebuild priority is "Normal." Rebuild will have less impact on host I/O access, but will take a longer time to complete. Changing the priority to a higher level will achieve a faster rebuild, but will significantly increase the host I/O response time. The default setting "Normal" is recommended.

5.9 Viewing and Editing Logical Drives and Drive Members

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives..," then press ENT.

```
View and Edit  
Logical Drives ↓
```

Press the up or down arrow keys to select the logical drive, then press ENT.

```
LG0 RAID5 DRV=3  
2012MB GD SB=1
```

Press the up or down arrow keys to select "View Drives..", then press ENT .

```
View Drives
```

Press the up or down arrow keys to scroll through the list of member drives.

```
Slot=* 1010MB  
LG=0 LN SEAGATE
```

5.9.1 Deleting a Logical Drive



WARNING!

Deleting a logical drive erases all data stored in it.

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives," then press ENT.

```
View and Edit  
Logical Drives ↓
```

Press the up or down arrow keys to select a logical drive, then press ENT.

```
LG0 RAID5 DRV=3  
2012MB GD SB=1
```

Use the up or down arrow keys to select "Delete Logical Drive," then press ENT .

```
Delete Logical  
Drive ..
```

Press ENT for two seconds to delete. The selected logical drive has now been deleted.

```
LG=0  
Not Defined ?
```

5.9.2

Deleting a Partition of a Logical Drive

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives..," then press ENT.

```
View and Edit
Logical Drives ↓
```

Press the up or down arrow keys to select a logical drive, then press ENT.

```
LG0 RAID5 DRV=3
2012MB GD SB=1
```

Press the up or down arrow keys to choose "Partition Logical Drive," then press ENT.

```
Partition
Logical Drive ..
```

The first partition's information will be shown on the LCD. Press the up or down arrow keys to browse through the existing partitions in the logical drive. Select a partition by pressing ENT for two seconds.

```
LG=0 Prt=0
200MB ?
```

Use the up or down arrow keys to change the number of the flashing digit to "0," then press ENT to move to the next digit. After changing all the digits, press ENT for two seconds.

```
LG=0 Prt=1
300MB ?
```

```
LG=0 Prt=2
600MB ?
```

The disk space of the deleted partition will be automatically allocated to the previous partition as diagrammed below. If partition 1 is deleted, its disk space will be added to partition 0.

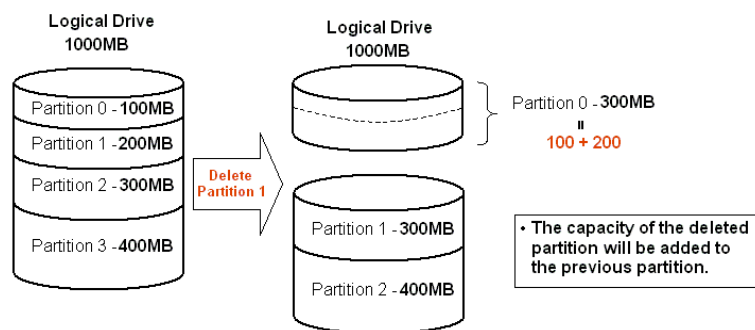


Fig. 5.1 Drive Space Allocated to the Previous Partition



WARNING!

Whenever there is a partition change, data will be erased. Prior to partition change, you have to remove its associated host LUN mappings. After the partition change, you also need to re-arrange the disk volumes from your host system OS.

5.9.3 Assigning a Name to a Logical Drive

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives..," then press ENT.

```
View and Edit
Logical Drives ↓
```

Press the up or down arrow keys to select a logical drive, then press ENT.

```
LG0 RAID5 DRV=3
  2012MB GD SB=1
```

Press the up or down arrow keys to select "Logical Drive Name," then press ENT.

```
Logical Drive
Name ..
```

Press the up or down arrow keys to change the character of the flashing cursor. Press ENT to move the cursor to the next space. The maximum number of characters for a logical drive name is 14.

```
Enter LD Name:
_
```

5.9.4 Rebuilding a Logical Drive (Manual Rebuild)

If you want the controller to auto-detect a replacement drive, make sure you have the following options set to enabled:

1. Periodic Drive Check Time
2. Periodic Auto-Detect Failure Drive Swap Check Time

These two configuration options can be found under "View and Edit Configuration Parameters" -> "Drive-Side Parameters".

A manual rebuild takes place after you replace the failed drive with a new drive that is known to be good.

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Logical Drives..", then press ENT.

```
View and Edit
Logical Drives ↓
```

Use the up or down arrow keys to select the logical drive that has a failed member, then press ENT.

```
LG0 RAID5 DRV=3
  2012MB FL SB=0
```

Use the up or down arrow keys to select "Rebuild Logical Drive," then press ENT.

```
Rebuild Logical
Drive ..
```

Press ENT for two seconds to start rebuilding the logical drive.

```
Rebuild Logical
Drive ?
```

The rebuilding progress will be displayed (as a percentage) on the LCD. When rebuild is already started or the logical drive is being rebuilt by a Local Spare Drive or Global Spare Drive, choose "Rebuild Progress" to see the rebuild progress.

```
LG0 RAID5 DRV=3
2012MB RB SB=0
```

```
Rebuilding 25%
Please Wait!
```

```
Rebuild Progress
..
```



NOTICE!

- The Rebuild function will appear only if a logical drive (with RAID level 1, 3 or 5) has a failed member.
- Carefully verify the physical location of a failed drive before replacement takes place. Removing the wrong drive may cause a logical drive to fail and data loss is unrecoverable.

5.9.5

Regenerating Logical Drive Parity

If no verifying method is applied to data writes, this function can be manually performed to ensure that parity errors can be mended.

From the Main Menu, press the up or down arrow keys to select "View and Edit Logical Drives."

```
View and Edit
Logical Drives
```

If you have more than one logical drive, use the up or down arrow keys to select the logical drive you would like to check the parity for, and then press ENT.

```
LG0 RAID5 DRV=3
4095MB GD SB=0
```

Press the up or down arrow keys to select "Regenerate Parity" and then press ENT.

```
Regenerate Parity
..
```

To stop the regeneration process, press ESC and enter the submenu to select "Abort Regenerate Parity".

```
Abort Regenerate
Parity ..
```



NOTICE!

If the Parity Regenerating process is stopped by a drive failure, the process cannot restart until the logical drive is rebuilt.

5.9.6

Media Scan

Media Scan is used to examine drives and detect the presence of bad blocks. If any data blocks have not been properly committed, data from those blocks are automatically recalculated, retrieved, rewritten or stored onto undamaged sectors. If bad blocks are encountered on yet another drive during the rebuild process, the block LBA (Logical Block Address) of those bad blocks will be listed. If rebuild is carried out in this situation, rebuild will continue with the unaffected sectors, salvaging a majority of the stored data.

From the Main Menu, press the up or down arrow keys to select "View and Edit Logical Drives".

```
View and Edit
Logical Drives
```

The first logical drive displays. If you have more than one logical drive, use the up or down keys to select the logical drive you want to scan, and then press ENT.

```
LG0 RAID5 DRV=3
4095MB GD SB=0
```

Press the up or down arrow keys to select "Media Scan" and then press ENT.

```
Media Scan
..
```

Press ENT again to display the first configuration option, "Priority." Press ENT on it and use arrow keys to select an option. Press ENT to confirm the change on priority level.

```
Priority
Normal ..
```

```
Priority
To High ?
```

Use arrow keys to move one level down to another option, "Iteration Count". This option determines how many times the scan is performed on the logical drive. If set to "Continuous," the scan will run in the background continuously until it is stopped by user.

```
Iteration Count
Single ..
```

```
Iteration Count
to Continuous ?
```

If Media Scan continuously runs in the background, considerable system resources will be consumed.

Press ENT on your option to confirm.

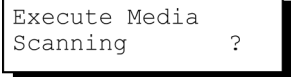
Press ENT for two seconds to display the confirm message, then press ENT to start scanning the array.

```
Execute Media
Scanning ?
```

5.9.7

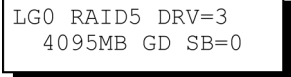
Write Policy

From the Main Menu, use the up or down arrow keys to select "View and Edit Logical Drives".



```
Execute Media
Scanning      ?
```

The first logical drive displays. If you have more than one logical drive, use the up or down keys to select the logical drive you want to change the write policy of; and then press ENT.




```
LG0 RAID5 DRV=3
4095MB GD SB=0
```

Use arrow keys to select "Write Policy" and then press ENT.



```
Write Policy
..
```

The Write-Back cache setting is configurable on a per array basis. Setting to the default value means the array setting is coordinated with the controller's general setting. The controller's general setting option can be found in "View and Edit Config Params" -> "Caching Parameters" -> "Write-Back Cache". Note that cached writes are lost if a power failure occurs unless cached data has been duplicated to a partner controller and a battery is supporting cache memory.



```
Write Policy
Write-Back
```


5.10 Viewing and Editing Host LUNs

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Host Luns", then press ENT.

```
View and Edit
Host Luns  ↑
```

Press the up or down arrow keys to select a host ID, then press ENT to proceed.

```
CH0 ID0
..
```

Use the up or down arrow keys to browse through the LUN number and its LUN mapping information.

```
CH0 ID0 LUN0
Mapto LG0 PRT0
```

Press ENT on the LUN you wish to delete.

```
Delete CH0 ID0
LUN=00 Mapping ?
```

Press ENT for two seconds to confirm deletion. The deleted LUN has now been unmapped.

```
CH0 ID0 LUN0
Not Mapped
```

5.11 Viewing and Editing Drives

Deleting Spare Drive (Global / Local Spare Drive)

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press ENT.

```
View and Edit
Drives  ↑
```

Drive information will be displayed on the LCD. Press the up or down arrow keys to select the spare drive you wish to delete, then press ENT.

```
Slot=* 1010MB
GlobalSB WDC
```

Press the up or down arrow keys to select "Delete Spare Drive," then press ENT to continue.

```
Delete Spare
Drive ..
```

Press ENT for two seconds to delete the spare drive.

```
Delete Spare
Drive Successful
```

5.12 Viewing and Editing Channels

5.12.1 Viewing IDs

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Channels," then press ENT.

```
View and Edit
Channels  ↑
```

Channel information will be displayed. Press ENT on the host channel you wish the ID changed.

```
CH0=Host ID=0
SXF=160M
```

Press the up or down arrow keys to select "Set SCSI Channel ID," then press ENT.

```
Set SCSI Channel
ID ..
```

Use the up or down arrow keys to browse through the existing host IDs. Press ENT on any ID combination to continue.

```
CHL=0 ID=0
..
```

5.12.2 Adding a Channel ID

Press ENT on a host channel, on "Set Channel ID", and then on an existing ID. Use the up or down arrow keys to choose "Add Channel SCSI ID", then press ENT.

```
Add Channel SCSI
ID
```

Use the up or down arrow keys to choose the ID you wish to add, then press ENT for two seconds to complete the process.

```
Add CHL=0 ID=0
?
```

A prompt will remind you to reset the subsystem for the configuration change to take effect. You may press ENT to reset the subsystem immediately or you may press ESC to continue adding other host IDs and reset the subsystem later.

```
Change Setting
Do Reset Ctlr ?
```

5.12.3

Deleting a Channel ID

Press ENT for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Channels," then press ENT.

```
View and Edit
Channels      ↕
```

The first host channel should appear. Press ENT to select a host channel.

```
CHL=Host ID=0
SXF=160M     ..
```

Press ENT on "Set SCSI Channel ID.."

```
Set SCSI Channel
ID              ..
```

A list of host channel and host ID combinations will appear. Use the up or down arrow keys to select an ID you wish to remove. Press ENT to select a channel ID combination.

```
CHL=0 ID=0
          ..
```

```
CHL=0 ID=0
          ..
```

You will then be prompted by the "Add Channel SCSI ID" option. Press the down arrow key to proceed.

```
Add Channel SCSI
ID
```

The "Delete Channel SCSI ID" option will appear. Press ENT to display the confirm box. Press ENT for two seconds to remove the ID.

```
Delete Channel
SCSI ID      ..
```

```
Delete ID=0
          ?
```

A prompt will remind you to reset the subsystem for the configuration change to take effect. You may press ENT to reset the subsystem immediately or you may press ESC to continue adding other host IDs and reset the subsystem later.

```
Change Setting
Do Reset Ctlr ?
```

NOTICE!

- Every time you make changes to channel IDs, you must reset the subsystem for the configuration to take effect.
- For IDs reserved in different subsystem configurations, please refer to the Installation Manual that came with your system. For subsystems featuring back-end PCB connection, firmware can detect their board types and automatically apply the preset IDs. There is no need to set IDs for these models.



5.12.4

Data Rate

This option is available in the configuration menu of SCSI host channel and the drive channel configuration menus of SATA-based subsystems. Default is "AUTO" and should work fine with most drives. Changing this setting is not recommended unless some particular bus signal issues occur.

Most SATA/ATA-based systems connect only one drive per SATA/ATA channel (4 for multi-lane with SATA-II). This helps to avoid a single drive failure from affecting other drives. The maximum mechanical performance of today's drives can reach around 30MB/second (sustained read). This is still far below the bandwidth of a drive channel bus. Setting the SATA/ATA bus speed to a lower value can get around some problems, but will not become a bottleneck to system performance.

Note that the SATA/ATA speed is the maximum transfer rate of the SATA/ATA bus in that mode. It does not mean the drive can actually carry out that amount of sustained read/write performance. For the performance of each drive model, please refer to the documentation provided by drive manufacturer.

Host Channel:

Select from Main Menu, "View and Edit Channels," and then a host channel you wish to change its data rate. Press ENT on the channel and use the arrow keys to find the "Set Transfer Speed" option. Press ENT on the Transfer Speed option to display "Chl=X Clk=160 Change to=160?", where "X" stands for the channel number.

```
C0=Host PID=0
SID=NA SXF=160M
```

```
Set Transfer
Speed ..
```

```
Chl=x Clk=160
Change to=160 ?
```

Use your arrow keys to display a different data rate (160MHz to Async). Press ENT to confirm a selection.

Drive Channel:

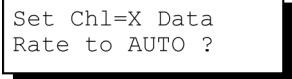
Select from Main Menu, "View and Edit Channels," and then a drive channel you wish to change its data rate. Press ENT on the channel and use the arrow keys to find the "Data Rate" option. Press ENT on the Data Rate option to display "Set Chl=X Data Rate To AUTO?", where "X" stands for the channel number.

```
CHX=Drive SATA
SXF=AUTO
```

```
View Chip
Information ..
```

```
Data Rate
..
```

Use your arrow keys to display a data rate value which ranges from 33 to 150MB/s (SATA drive channels). Press ENT to confirm a selection.

A rectangular LCD screen with a black border. The screen displays two lines of text in a monospaced font. The first line is "Set Chl=X Data" and the second line is "Rate to AUTO ?".

Set Chl=X Data
Rate to AUTO ?

5.13 System Functions

Choose "System Functions" in the Main Menu, then press ENT. Press the up or down arrow keys to select a submenu, then press ENT.

5.13.1 Mute Beeper

When the controller's beeper has been activated, choose "Mute Beeper," then press ENT to turn the beeper off temporarily for the current event. The beeper will still activate on the next event.



Mute Beeper ..

5.13.2 Password

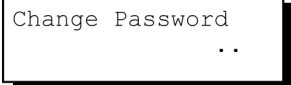
Use the controller's password to protect the system from unauthorized entry. Once the controller's password is set, regardless of whether the front panel, the RS-232C terminal interface or the RAIDWatch Manager is used, the user can only configure and monitor the RAID controller by providing the correct password.

NOTICE!

- The controller requests a password whenever a user is entering the main menu from the initial screen or a configuration change is made. If the controller is going to be left unattended, the "Password Validation Timeout" should be set to "Always Check."
- The controller password and controller name share a 32-character space. The maximum number of characters for a controller password is 32. If 31 characters are used for a controller name, there will be only one character left for the controller password and vice versa. Since 3.42J09, later firmware revisions will support a 32-character space.

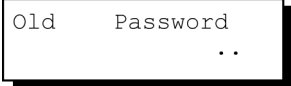
Change Password

To set or change the controller password, press the up or down arrow keys to select "Change Password," then press ENT.



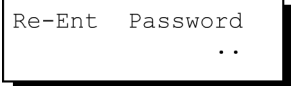
Change Password ..

If the password has previously been set, the controller will ask for the old password first. If password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

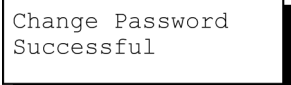


Old Password ..

Press the up or down arrow keys to select a character, then press ENT to move to the next space. After entering all the characters (alphabetic or numeric), press ENT for two seconds to confirm. If the password is correct, or there is no preset password, it will ask for the new password. Enter the password again to confirm.



Re-Ent Password ..



Change Password
Successful

Disable Password

To disable or delete the password, press ENT on the first flashing digit for two seconds when requested to enter a new password. The existing password will be deleted. No password checking will occur when entering the Main Menu from the initial terminal screen or making configuration changes.

5.13.3

Reset Controller

To reset the controller without powering off the system, Press the up or down arrow keys to “Reset Controller,” then press ENT. Press ENT again for two seconds to confirm. The controller will now reset.

```
Reset This  
Controller  ..
```

```
Reset This  
Controller  ?
```

5.13.4

Shutdown Controller

Before powering off the controller, unwritten data may still reside in cache memory. Use the “Shutdown Controller” function to flush the cache content. Press the up or down arrow keys to “Shutdown Controller,” then press ENT. Press ENT again for two seconds to confirm.

```
Shutdown This  
Controller  ?
```

```
Shutdown This  
Controller  ..
```

The controller will now flush the cache memory. Press ENT for two seconds to confirm and to reset or power off the subsystem.

```
ShutdownComplete  
Reset Ctlr?
```

5.13.5

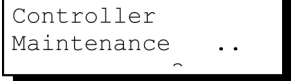
Controller Maintenance

For Controller Maintenance functions, see *Section 16 Event Messages*.

5.13.6 Saving NVRAM to Disks

You can choose to backup your controller-dependent configuration information to disks. We strongly recommend using this function to save the configuration profile whenever a configuration change is made. The information will be distributed to every logical drive in the RAID system. If using the RAIDWatch manager, you can save your configuration details as a file to a computer system drive.

A RAID configuration of drives must exist for the controller to write NVRAM content onto it. From the Main Menu, choose "System Functions." Use arrow keys to scroll down and select "Controller Maintenance," "Save NVRAM to Disks," then press ENT. Press ENT for two seconds on the message prompt, "Save NVRAM to Disks?".



```
Controller
Maintenance  ..
```



```
Save NVRAM
To Disks    ?
```

A prompt will inform you that NVRAM information has been successfully saved.

5.13.7 Restore NVRAM from Disks

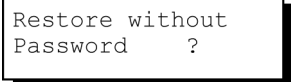
If you want to restore your NVRAM information that was previously saved onto the array, use this function to restore the configuration setting.

From the Main Menu, choose "System Functions." Use arrow keys to scroll down and select "Controller Maintenance," "Restore NVRAM from Disks..," and then press ENT. Press ENT for two seconds to confirm.

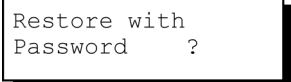


```
Restore NVRAM
from Disks    ?
```

In case your previous password (reserved at the time you saved your NVRAM configuration contents) is different from your current password, you are provided with the options whether to restore the password you previously saved with your configuration profile.



```
Restore without
Password      ?
```

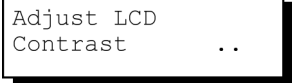


```
Restore with
Password     ?
```

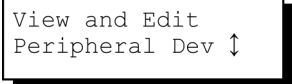
A prompt will inform you that the controller NVRAM data has been successfully restored from disks.

5.13.8 Adjust LCD Contrast

The controller LCD contrast is set at the factory to a level that should be generally acceptable. The controller is equipped with an LCD contrast adjustment circuit in case the factory-preset level needs to be adjusted either via the RS-232 terminal emulation menus or using the LCD keypad panel.



```
Adjust LCD  
Contrast ..
```



```
View and Edit  
Peripheral Dev ↑
```

From the main menu, choose “View and Edit Peripheral Dev.” Press ENT on it, press arrow keys to scroll down, and select “Adjust LCD Contrast,” press ENT to proceed, and then use the arrow keys to find an optimal setting. Press ESC to return to the previous menu.



```
LCD Contrast :  
■■■■■■■■-----
```

5.14 Controller Parameters

5.14.1 Controller Name

Select "View and Edit Config Parm's" from the Main Menu. Choose "View and Edit Configuration Parameters," "Controller Parameters," then press ENT. The current name will be displayed. Press ENT for two seconds and enter the new controller name by using the up or down arrow keys. Press ENT to move to another character and then press ENT for two seconds on the last digit of the controller name to complete the process.

```
Controller
Name:
```

```
Enter Ctlr Name:
□
```

5.14.2 LCD Title Display Controller Name

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press ENT. Use the up or down arrow keys to choose to display the embedded controller logo or any given name on the LCD initial screen.

```
LCD Title Disp -
Controller Logo?
```

```
LCD Title Disp -
Controller Name?
```

5.14.3 Password Validation Timeout

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press ENT. Select "Password Validation Timeout," and press ENT. Press the up or down arrow keys to choose to enable a validation timeout from one to five minutes to Always Check. The Always Check timeout will invalidate any configuration change without entering the correct password.

```
PasswdValidation
Timeout-5 mins..
```

5.14.4 Controller Unique Identifier

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press ENT. Press the up or down arrow keys to select "Ctlr Unique ID-," then press ENT. Enter any hex number between "0" and "FFFF" and press ENT to proceed.

```
Ctlr Unique (Hex)-
ID-
```

Enter a unique ID for any RAID controller in a single or dual-controller configuration. The unique ID is recognized by the controller as the following:

1. The unique ID is combined to generate a unique WWN node name for controllers or RAID systems using Fibre Channel host ports. The unique node name helps prevent host computers from mis-addressing the storage system during a controller failback/failover process.
2. MAC addresses for the controller's Ethernet port that should be taken over by a surviving controller in the event of controller failure.

5.14.5 Controller Date and Time

This submenu is only available for controllers or subsystems that come with a real-time clock on board.

Time Zone

Choose “View and Edit Configuration Parameters,” “Controller Parameters,” then press ENT. Press the up or down arrow keys to scroll down and select “Set Controller Date and Time”, then press ENT.

```
View and Edit  
Config Params
```

```
Controller  
Parameters ..
```

```
Set Controller  
Date and Time ..
```

The controller uses GMT (Greenwich Mean Time), a 24-hours clock. To change the clock to your local time zone, enter the hours later than the Greenwich mean time after a plus (+) sign. For example, enter “+9” for Japanese time zone. Choose “Time Zone” by pressing ENT.

```
Time Zone ..
```

Use the down key to enter the plus sign and the up key to enter numeric representatives.

```
GMT +08:00  
GMT -
```

Date and Time

Use your arrow keys to scroll down and select “Date and Time” by pressing ENT.

```
Date and Time ..
```

Use the arrow keys to select and enter the numeric representatives in the following order: month, day, hour, minute, and the year.

```
[MMDDhhmm[YYYY]]
```

5.15 Drive Utilities

From the "View and Edit Drives" menu, select the drive that the utility is to be performed on; then press ENT. Select "Drive Utilities;" then press ENT. Choose "Read/Write Test".

```
View and Edit
Drives
```

```
Slot=* 1010MB
NEW DRV SEAGATE
```

These options are not available for drives already configured in a logical configuration, and can only be performed before a reserved space is created on a drive.

```
Drive
Utilities  ..
```

```
Drive Read/Write
Test  ..
```

5.15.1 Drive Read/Write Test

From the "View and Edit Drives" menu, select a new or used drive that the utility is to be performed on; then press ENT. Select "SCSI Drive Utilities;" then press ENT. Choose "Read/Write Test" and press ENT.

```
Drive Read/Write
Test  ..
```

Press the up or down arrow keys to select and choose to enable/disable the following options:

1. Abort When Error Occurs
2. Drive Test for - Read Only/Read and Write.
3. Execute Drive Testing

```
Abort When Error
Occur-Enabled ....
```

```
Drive Test for
Read and Write..
```

```
Execute Drive
Testing  ..
```

When finished with configuration, select "Execute Drive Testing" and press ENT to proceed. The Read/Write test progress will be indicated as a percentage.

```
Drv Testing 23%
Please Wait !
```

You may press ESC and select "Read/Write Test" later and press or ▲ to select to "View Read/Write Testing Progress" or to "List Current Bad Block Table." If you want to stop testing the drive, select "Abort Drive Testing" and press ENT to proceed.

View Read/Write
Test Progress ..

List Current
Bad Block Table.

Abort Read/Write
Testing ..

6 Terminal Screen Messages

6.1 The Initial Screen




Fig. 6.1 The RC-232C Initial Screen

Cursor Bar:	Move the cursor bar to a desired item, then press [ENTER] to select
Subsystem Name:	Identifies the type of controller/subsystem or a preset name
Transfer Rate Indicator:	Indicates the current data transfer rate
Gauge Range:	Use + (Shift+"") or - keys to change the gauge range in order to view the transfer rate indicator
Cache Status:	Indicates current cache status
Write Policy:	Indicates current write-caching policy
Date & Time:	Current system date and time, generated by controller real-time clock
PC Graphic (ANSI Mode):	Enters the Main Menu and operates in ANSI mode
Terminal (VT-100 Mode):	Enters the Main Menu and operates in VT-100 mode
PC Graphic (ANSI+Color Mode):	Enters the Main Menu and operates in ANSI color mode
Show Transfer Rate+Show Cache Status:	Press [ENTER] on this item to show the cache status and transfer rate
Ongoing Processing:	e#: logical drive # is being expanded i#: logical drive # is being initialized R#: logical drive # is being rebuilt P#: logical drive # Parity Regeneration completion ratio S#: logical drive # Media Scan completion ratio For more details, please refer to <i>Section 6.4 Logical Drive Status</i> .

6.2 Main Menu

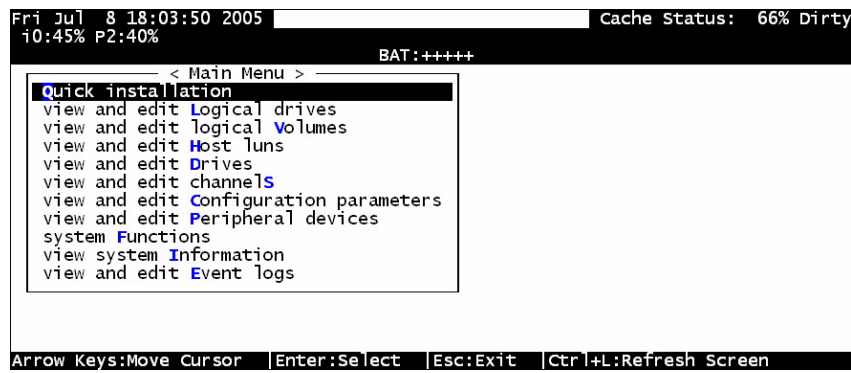


Fig. 6.2 The RC-232C Main Menu

Use the arrow keys to move the cursor bar through the menu items, then press [ENTER] to choose a menu, or [ESC] to return to the previous menu/screen.

In a subsystem or controller head where battery status can be detected, battery status will be displayed at the top center. Status will be stated as Good, Bad, several “+” (plus) signs (VT-100 mode), or color blocks (ANSI mode) will be used to indicate battery charge. A battery fully-charged will be indicated by five plus signs or color blocks.

When initializing or scanning an array, the controller displays progress percentage on the upper left corner of the configuration screen. An “i” indicates array initialization. An “s” stands for scanning process. The number(s) next to them indicate the logical drive number (e.g., logical drive 0).

6.3 Quick Installation

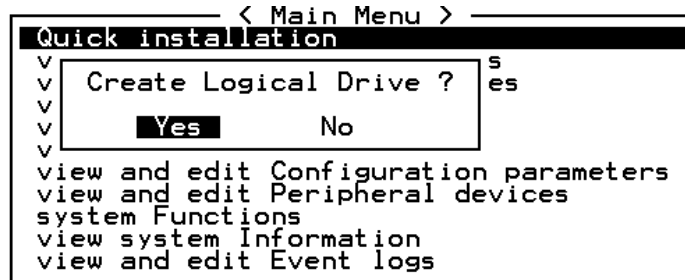


Fig. 6.3 Quick Installation Page

Type Q or use the ↑↓ keys to select "Quick installation", then press [ENTER]. Choose Yes to create a logical drive.

All possible RAID levels will be displayed. Use the ↑↓ keys to select a RAID level, then press [ENTER]. The assigned spare drive will be a Local Spare Drive, not a Global Spare Drive.

All available disk drives in the enclosure will be included in one logical drive. The subsystem will start initialization and automatically map the logical drive to LUN 0 of the first host channel available.

6.4 Logical Drive Status

Cache Status: Clean
Write Cache: Enable

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0	62EC6758	NA	RAID5	5488	GOOD					6	3	0	0	
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 6.4 Logical Drive Status Screen



NOTICE!

A logical drive in a single-controller subsystem is always managed by one controller, and the “P” or “S” indicator will not appear.

LG	Logical Drive number
LV	The Logical volume to which this logical drive belongs
ID	Controller-generated unique ID
RAID	RAID level
SIZE (MB)	Capacity of the Logical Drive
Status 1	Logical Drive Status – Column 1
GOOD	The logical drive is in good condition
DRV FAILED	A drive member failed in the logical drive
CREATING	Logical drive is being initiated
DRV ABSENT	An empty drive tray
INCOMPLETE	Two or more drives failed in the logical drive
INVALID	The logical drive was created but has not been fully initialized when another version of firmware is being loaded. After the subsystem resets, the array status should return to normal.
FATAL FAIL	Two member drives failed at the same time, the array is inaccessible
DRV MISS	A member drive is missing; could result from insecure installation
REBUILDING	The logical drive is being rebuilt
Status 2	Logical Drive Status – Column 2
I	Initializing drives
A	Adding drive(s)
E	Expanding logical drive
H	Add drive operation on hold
Status 3	Logical Drive Status – Column 3
R	Rebuilding the logical drive
P	Regenerating array parity

Column O		Logical Drive Status – Stripe size		
	N/A	Default		
	2	4KB	7	128KB
	3	8KB	8	256KB
	4	16KB	9	512KB
	5	32KB	A	1024KB
	6	64KB		
Column C		Logical Drive Status – Write Policy setting		
	B	Write-back		
	T	Write-through		
#LN		Total drive members in the logical drive		
#SB		Standby drives available for the logical drive. This includes all the spare drives (local spare, global spare) available for the specific logical drive		
#FL		Number of Failed member(s) in the logical drive		
Name		Logical drive name (user configurable)		

6.5 Logical Volume Status

					Cache Status: Clean	
Q	LV	ID	Size(MB)	#LD	s es parameters vices	
v	0	466C5C8D	600000	1		
v	1					
v	2					
v	3					
s	4					
v	5					
v	6					
	7					
Arrow Keys: Move Cursor Enter: Select Esc: Exit Ctrl+L: Refresh Screen						

Fig. 6.5 Logical Volume Status Screen



NOTICE!

A logical volume in a single-controller subsystem is always managed by one controller, and the “P” or “S” indicator will not appear.

LV	Logical Volume number.
ID	Logical Volume ID number (controller randomly generated)
Size (MB)	Capacity of the Logical Volume
#LD	The number of Logical Drive(s) included in this Logical Volume

6.6 Drive Status

Fri Sep 12 18:07:09 2003				Cache Status: Clean			
< Main Menu >				BAT: █			
Quick installation							
view and edit logical drives							
view and edit logical volumes							
view and edit host lun							
view and edit drives							
view and edit channels							
view and edit configuration							
view and edit peripheral							
system functions							
view system information							
view and edit event log							
Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID		
5	238469	150MB	GLOBAL	STAND-BY	WDC	WD2500JD-50FYB0	
6				ABSENT			
7				ABSENT			
8	238469	150MB	NONE	USED DRV	WDC	WD2500JD-00FYB0	
9	238469	150MB	1	ON-LINE	WDC	WD2500JD-50FYB0	
10	238469	150MB	1	ON-LINE	WDC	WD2500JD-50FYB0	
11	238469	150MB	1	ON-LINE	WDC	WD2500JD-50FYB0	

Fig. 6.6 Drive Status Screen

Slot	Slot number of the disk drive	
Size (MB)	Drive capacity	
Speed	XXMB Maximum transfer rate of the drive channel interface	
LG_DRV	X	The disk drive is a member of logical drive "X." If the Status column shows "STAND-BY", the drive is a Local Spare belonging to logical drive "X."
Status	Global	The disk drive is a Global Spare Drive
	INITING	Proceeding with array initialization
	ON-LINE	The drive is in good condition
	REBUILD	Proceeding with array Rebuild process
	STAND-BY	Local Spare Drive or Global Spare Drive. The Local Spare Drive's LG_DRV column will show the logical drive number. The Global Spare Drive's LG_DRV column will show "Global".
	NEW DRV	A new drive has not been included in any logical drive or configured as a spare drive
	USED DRV	An used drive that is not a member of any logical drive or configured as a spare
	FRMT DRV	Formatted drive (drive formatted with a reserved section)
	BAD	Failed drive
	ABSENT	Drive does not exist
	MISSING	Drive once existed, but is missing now
	SB-MISS	Spare drive missing
Vendor and Product ID	The vendor and product model information of the drive	

6.7

Channel's Status

Thu Sep 1 16:23:13 2005

Cache Status: Clean

BAT:NONE

Q	Ch1	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	0	160.0MHz	Wide	L	On	Async	Narrow	
1	Host	0	160.0MHz	Wide	L	On	Async	Narrow	
2	Drive		AUTO	SATA				SATA	
3	Drive		AUTO	SATA				SATA	
4	Drive		AUTO	SATA				SATA	
5	Drive		AUTO	SATA				SATA	
6	Drive		AUTO	SATA				SATA	
7	Drive		AUTO	SATA			1.33GHZ	SATA	

Arrow Keys: Move Cursor

Enter: Select

Esc: Exit

Ctrl+L: Refresh Screen

Fig. 6.7 Channel Status Screen

Chl	Channel number	
Mode	Channel mode	
	Host	Host Channel mode
	Drive	Drive Channel mode
ID	IDs managed by the Controller	
	*	Multiple IDs were applied (Host Channel mode only)
	(ID number)	Host Channel: Specific IDs managed by the Controller for host LUN mapping
		Drive Channel: Specific ID reserved for the channel processor on the Controller
AUTO	Channel bus data rate set to auto speed negotiation	
DefSynClk	Default bus synchronous clock:	
	???.?M	The default setting of the channel is ??? MHz in Synchronous mode.
	Async.	The default setting of the channel is Asynchronous mode.
DefWid	Default bus width:	
	Width	SCSI transfer protocol
	SATA	As regulated by SATA specifications
S	Signal:	
	L	SCSI-320
Term	Terminator Status:	
	On	Terminator is enabled.
	Off	Terminator is disabled.
	Diff	The channel is a Differential channel. The terminator can only be installed/removed physically.
	N/A	Non-SCSI bus

CurSynClk	Current bus synchronous clock:	
	???.?GHz	The default setting of the channel bus is ???.? GHz
	Async.	The default setting of the channel bus is Asynchronous mode.
	(empty)	The default bus synchronous clock has changed. Reset the controller for the changes to take effect.
CurWid	Current Bus Width:	
	Narrow	SCSI transfer protocol
	SATA	As regulated by SATA specifications

6.8

Controller Voltage and Temperature

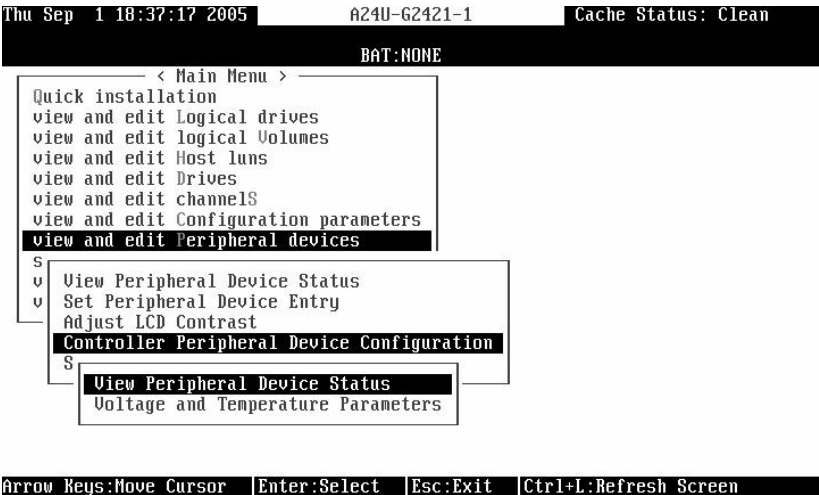


Fig. 6.8 Controller Voltage and Temperature Monitoring Page

Choose from Main Menu “View and Edit Peripheral Devices,” and press [ENTER]. From the submenu, choose “Controller Peripheral Device Configuration,” “View Peripheral Device Status”, then press [ENTER].

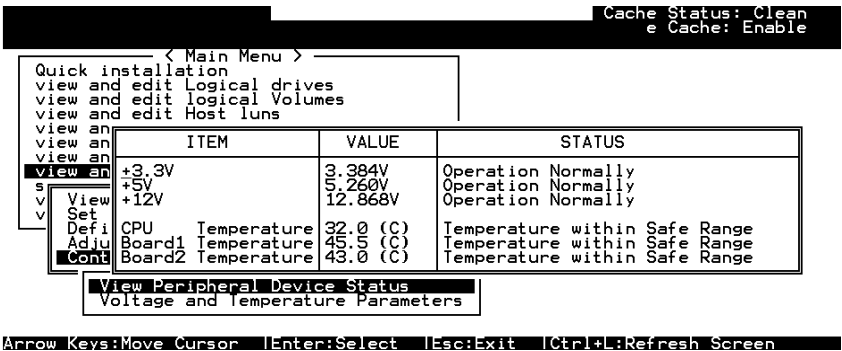


Fig. 6.9 Controller Voltage and Temperature Status Screen

The current status of voltage and temperature detected by the controller will be displayed on-screen and will be stated as normal, out of order, within or within the safety range.

6.9 Viewing Event Logs on the Screen

When errors occur, you may want to trace the records to see what has happened to your system. The controller's event log management records all events starting from the time when the system is powered on, recording up to 1,000 events. Powering off or resetting the controller will cause an automatic deletion of all the recorded event logs. To view the events log on-screen, from the Main Menu "View and Edit Event Logs" by pressing [ENTER].

```

Sun Feb 17 23:08:08 2002                               Cache Status: Clean
BAT:++++
Event Logs
On-Line Initialization of Logical Drive 0 Completed
<Sun Feb 17 22:48:22 2002>                               <P>
LG:0 Logical Drive NOTICE: Starting On-Line Initialization
<Sun Feb 17 22:47:49 2002>                               <P>
Creation of Logical Drive 0 Completed
<Sun Feb 17 22:47:49 2002>                               <P>
LG:0 Logical Drive NOTICE: Starting Creation
<Sun Feb 17 22:47:44 2002>                               <P>
Controller BBU Not Fully Charged !
<Sun Feb 17 22:33:24 2002>                               <S>
Controller Initialization Completed
<Sun Feb 17 22:33:27 2002>                               <S>
Controller NOTICE: NVRAM Factory Defaults Restored
<Sun Feb 17 22:32:22 2002>                               <S>
Controller Initialization Completed
<Sun Feb 17 22:33:22 2002>                               <P>
Arrow Keys:Move Cursor | Enter:Select | Esc:Exit | Ctrl+L:Refresh Screen

```

Fig. 6.10 Event Logs Screen

To clear the saved event logs, scroll the cursor down to select an event and press [ENTER] to delete the event and the events below.

Clear Below Event Logs ?

Yes **No**

Fig. 6.11 Clear Event Logs Dialogue Box

Choose Yes to clear the recorded event logs.

7 Terminal Operation

7.1 Power on RAID Enclosure

Hardware installation should be completed before powering on your RAID enclosure. The sub-system and disk drives must be configured and properly initialized before the host computer can access the storage capacity. The configuration and administration utility resides in the controller's firmware.

Open the initial terminal screen: use the arrow keys to move the cursor bar through the menu items, then press [ENTER] to choose the terminal emulation mode, and [ESC] to return to the previous menu/screen.




Fig. 7.1 RS-232C Initial Screen



NOTICE!

All figures in this chapter are showing examples using the management hyper terminal screen.

7.2 Caching Parameters

7.2.1 Optimization Modes

Mass storage applications can be roughly divided into two categories—database and video/imaging—according to their read/write characteristics. To optimize system operation, there are two embedded optimization modes with system behaviors pre-adjusted to different read/write parameters: “Optimization for Random I/O” and the “Optimization for Sequential I/O.”

Limitations

The following are the limitations on the use of optimization modes.

1. The system default is Optimization for Sequential I/O.
2. You can select the stripe size of each array (logical drive) during the initial configuration. However, changing stripe size is only recommended for experienced engineers who have tested the effects of tuning stripe sizes for different applications.
3. The array stripe size can only be changed during the initial configuration process.
4. Carefully consider the outcome before choosing an optimization mode. Once the controller optimization mode is applied, access to different arrays in a RAID system will follow the same optimized pattern. You can only change the optimization mode after recreating the arrays.

Database and Transaction-based Applications

These kinds of applications usually include SQL server, Oracle server, Informix, or other database services that keep the size of each transaction down to a minimum, so that I/Os can be rapidly processed. Due to their transaction-based nature, these applications do not read or write a bunch of data in a sequential order—access to data occurs randomly. The transaction size usually ranges from 2K to 4K. Transaction performance is measured in “I/Os per second” or “IOPS.”

Video Recording/Playback and Imaging Applications

These kinds of applications usually include video playback, video post-production editing, or applications of a similar nature that tend to read or write large files to and from storage in a sequential order. The size of each I/O can be 128K, 256K, 512K, or up to 1MB. The efficiency of these applications is measured in MB per second (MBps).

When an array works with applications such as video or image-oriented applications, the application reads/writes from the drive as large-block, sequential threads instead of small and randomly accessed files.

The controller optimization modes have read-ahead buffer and other Read/Write characteristics tuned to obtain the best performance for these two major application categories.

7.2.2 Optimization Mode and Stripe Size

Each controller optimization mode has preset values for the stripe size of arrays created in different RAID levels. If you want a different optimization mode for a configured array, you must backup or move the stored data and recreate the arrays following the steps below:

1. Stop host I/O access.
2. Move or backup all of your stored data in the subsystem.
3. Change the optimization mode.
4. Reset the subsystem.
5. Recreate the array(s). Once the array(s) are created, stripe size cannot be changed.

Listed below are the default stripe sizes implemented with different optimization modes and RAID levels. These values should be sufficient for most applications.

	Stripe Size: Opt. for Sequential I/O	Stripe Size: Opt. for Random I/O
RAID0	128KB	32KB
RAID1	128KB	32KB
RAID3	16KB	4KB
RAID5	128KB	32KB
NRAID	128KB	32KB

Table 7.1 RAID Levels, Optimization Modes, and Stripe Sizes

7.2.3

Optimization for Random or Sequential I/Os

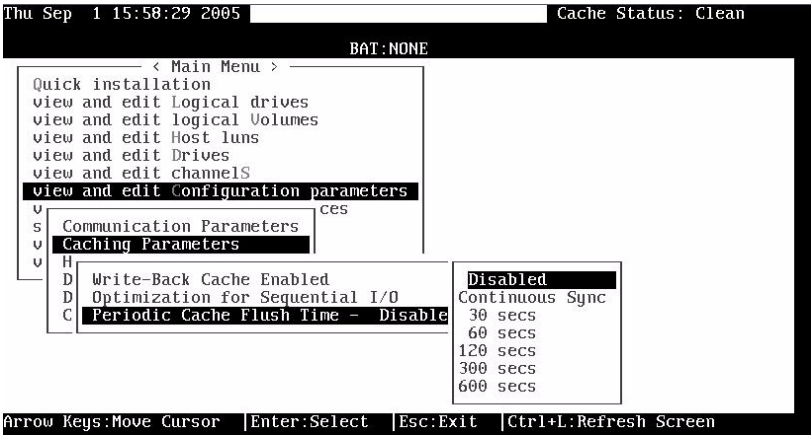


Fig. 7.2 Optimization for Random or Sequential I/Os

Page

Choose “Optimization for Random I/O” or “Optimization for Sequential I/O,” then press [ENTER]. The “Random” or “Sequential” dialog box will appear, depending on the option you have selected. Choose Yes in the dialog box that follows to confirm the setting. Press [ESC] to exit and the setting will take effect after the subsystem is restarted.



NOTICE!

There is no longer a 512GB threshold for array optimization modes. If an array is larger than 16TB, only the optimization mode for sequential I/Os can be applied. Since logical drives of this size may not be practical, there is actually no limitation on the optimization mode and array capacity.

7.2.4

Write-Back Cache Enable/Disable

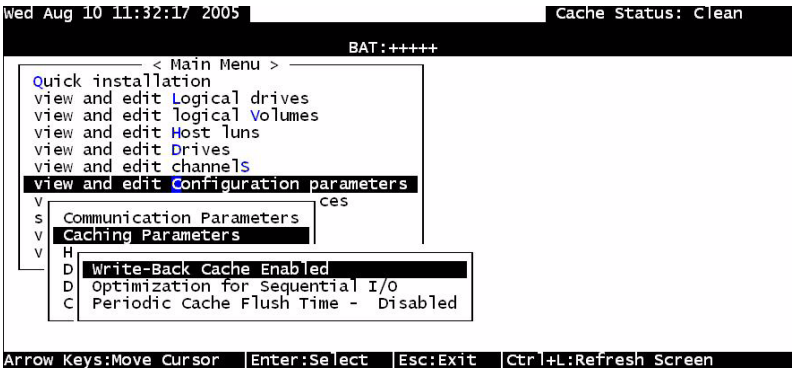


Fig. 7.3 Write-Back Cache Enable/Disable Page

As one of the submenus in “Caching Parameters,” this option controls the cached write policy. Choose *Caching Parameters*, then press [ENTER]. Select *Write-Back Cache*, then press [ENTER]. “Enabled” or “Disabled” will display the current setting with Write-back caching. Choose Yes in the dialog box that follows to confirm the setting.

Write-back caching can dramatically improve write performance by caching the unfinished writes in memory and letting them be committed to drives in a more efficient manner. In the event of power failure, a battery backup module can hold cached data for days.

7.2.5

Periodic Cache Flush

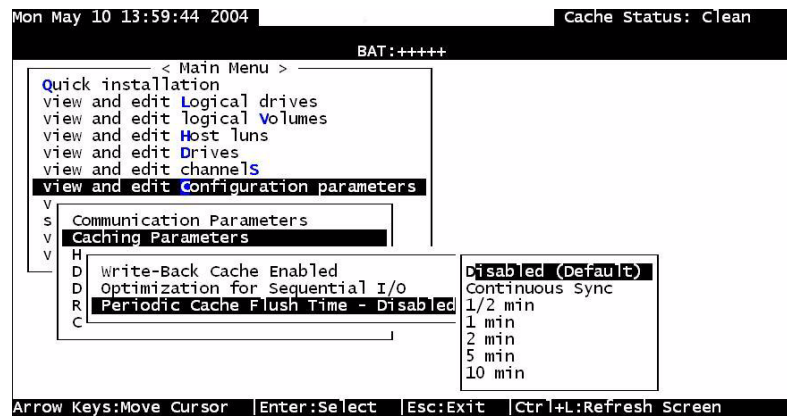


Fig. 7.4 Periodic Cache Flush Page

If Write-back caching is preferred for better performance yet data integrity is also a concern, e.g., no battery protection, the system can be configured to flush the cached writes at preset intervals.

Note that the “Continuous Sync” option holds data in cache for as long as necessary to complete a write operation and immediately commits it to hard drives if it does not come in a series of sequential write requests.

NOTICE!

- If the size of an array is larger than 16TB, only the optimization for sequential I/O can be applied. Since a logical drive of this size may not be practical, there is actually no limitation on the combination of optimization mode and array capacity.
- Every time you change the Caching Parameters you must reset the controller for the changes to take effect.
- The Adaptive Write Policy is applicable to subsystems working under normal conditions. If, for example, a drive fails in an array, the firmware automatically restores the array's original write policy.



7.3 Viewing the Connected Drives

Prior to configuring disk drives into a logical drive, it is necessary to understand the status of the physical drives in your enclosure.

Fri Jul 19 15:34:44 2002

Cache Status: Clean

BAT:+++++

Quick view view view view view view view view view	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID	
	1	239367	150MB	NONE	NEW DRV	Maxtor	7Y250M0
	2	239367	150MB	NONE	NEW DRV	Maxtor	7Y250M0
	3	239367	150MB	NONE	NEW DRV	Maxtor	7Y250M0
	4	239367	150MB	GLOBAL	STAND-BY	Maxtor	7Y250M0
	5	239367	150MB	0	ON-LINE	Maxtor	7Y250M0
	6	239367	150MB	1	ON-LINE	Maxtor	7Y250M0
	7	239367	150MB	1	ON-LINE	Maxtor	7Y250M0
	8	239367	133MB	0	ON-LINE	Maxtor	7Y250M0

Arrow Keys:Move Cursor | Enter:Select | Esc:Exit | Ctrl+L:Refresh Screen

Fig. 7.5 View and Edit Drives Page

Use the arrow keys to scroll down to “View and Edit Drives” to display information on all the physical drives installed.

Drives are listed in the “View and Edit Drives” table. Use the arrow keys to scroll the table. First examine whether there is any drive installed but not listed here. If a disk drive is installed but not listed, the drive may be defective or not installed correctly. If so, please contact your RAID supplier.

NOTICE!

- Drives of the same brand/model/capacity might not have the same block number.
- The basic read/write unit of a hard drive is block. If members of a logical drive have different block numbers (capacity), the smallest block number will be taken as the maximum capacity to be used in every drive. Therefore, use drives of the same capacity.
- You may assign a Spare Drive to a logical drive whose members have a block number equal or smaller than the Local/Global Spare Drive, but you should not do the reverse.



7.4 Creating a Logical Drive

Browse through the Main Menu and select *View and Edit Logical Drive*.

Thu Jul 21 17:18:22 2005 Cache Status: Clean

BAT:++++

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0			NONE											
Create Logical Drive ?														
Yes No														
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

Arrow Keys:Move Cursor | Enter:Select | Esc:Exit | Ctrl+L:Refresh Screen

Fig. 7.6 View and Edit Logical Drive Page

For the first logical drive on the RAID subsystem, simply choose the first logical drive entry, *LG 0*, and press [ENTER] to proceed. You may create as many as 16 logical drives or more using drives in a RAID subsystem or in a cascaded enclosure.

When prompted to “Create Logical Drive?,” select **Yes** and press [ENTER] to proceed.

Create Logical Drive ?	
Yes	No

7.4.1 Choosing a RAID Level

A pull-down list of supported RAID levels will appear. Choose a RAID level for this logical drive.

7.4.2 Choosing Member Drives

Choose your member drive(s) from the list of available physical drives. Tag the drives for inclusion by positioning the cursor bar on the drive and then pressing [ENTER]. A star sign “*” will appear on the selected physical drive(s). To deselect the drive, press [ENTER] again on the selected drive and the “*” will disappear. Use the same method to select more drives.

Thu Jul 21 17:24:13 2005 Cache Status: Clean

BAT:++++ 4 of 4 Selected

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	RAID 5	RAID 6	#FL	NAME
0													
Slot Size(MB) Speed LG_DRV Status Vendor and Product ID													
1													
2													
3													
4													
5													
6			NONE										
7			NONE										

Arrow Keys:Move Cursor | Enter:Select | Esc:Confirm | Ctrl+L:Refresh Screen

Fig. 7.7 Choose Member Drives Page

7.4.3

Logical Drive Preferences




Fig. 7.8 Logical Drive Preferences Screen

After all member drives have been selected, press [ESC] to continue with the next option. A list of array options is displayed.

Maximum Drive Capacity




Fig. 7.9 Maximum Drive Capacity Screen

As a rule, a logical drive should be composed of drives of the same capacity. A logical drive can only use the capacity of each drive up to the maximum capacity of the smallest member selected for the array.

Assign Spare Drives




Fig. 7.10 Assign Spare Drives Page

You can add a spare drive from the list of the unused drives. The spare chosen here is a Local spare and will automatically replace any failed drive. The controller will then rebuild data onto the replacement drive.

A logical drive composed in a non-redundancy RAID level (NRAID or RAID 0) does not support spare drive rebuild.

Disk Reserved Space

The reserved space is a small section of disk space formatted for storing array configuration and RAIDWatch program data. This item is for display only—you cannot change the size of the reserved space.

Write Policy




Fig. 7.11 Write Policy Sub-menu

This sub-menu allows you to set the caching mode for this specific logical drive. “Default” is a neutral value that is coordinated with the subsystem’s general caching mode setting shown in brackets in the Write Policy status.

Initialization Mode

```

Maximum Drive Capacity : 238214MB
Assign Spare Drives
Disk Reserved Space: 256 MB
Write Policy: Default(Write-Back)
Initialize Mode: On-Line
S Off-line Initialize for Create Logical Drive ?
7 Yes No

```

Fig. 7.12 Initialization Mode Sub-menu

This sub-menu allows you to see if the logical drive is immediately available. If the online (default) mode is used, data can be written onto it immediately and you may continue with array configuration, e.g., including the array into a logical volume, before the array’s initialization process is completed.

Stripe Size

```

M
A
D
D
L
W
I
Default (128K Bytes)
4K Bytes
8K Bytes
16K Bytes
32K Bytes
64K Bytes
128K Bytes
256K Bytes
512K Bytes
1024K Bytes
239112MB
MB
te-Back)
Stripe Size: Default (128K Bytes)

```

Fig. 7.13 Stripe Size Option Box

This option should only be changed by experienced engineers. Setting to an incongruous value can severely drag performance, therefore, this option should only be changed when you can be sure of the performance gains it might bring you.

The default value is determined by the combination of the controller Optimization Mode setting and the RAID level used for the array.

Press [ESC] to continue when all the preferences have been set.

```

Raid Level      : RAID 6
Online Drives   : 4
Maximum Drive Capacity : 239112 MB
Disk Reserved Space : 256 MB
Spare Drives    : 1
Logical Drive Assignment: Primary Controller
Write Policy     : Default(write-Back)
Initialize Mode  : On-Line
Stripe Size     : Default (128K Bytes)

Create Logical Drive ?
Yes No

```

Fig. 7.14 Logical Drive Preferences Screen


A confirm box will appear on the screen. Verify all information in the box before choosing Yes to confirm and proceed.

```


Notification
LG:0 Logical Drive NOTICE: Starting Creation

```

If the online initialization mode is applied, the logical drive will first be created and the controller will find an appropriate time to initialize the array.




The completion of array creation is indicated by the message prompt above.



A controller event will then prompt to indicate that the logical drive initialization has begun. Press [ESC] to cancel the “Notification” prompt, and a progress indicator will display on the screen as a percentage bar.

While the array initialization runs in the background, you can start using the array or continue configuring your RAID subsystem.



When a fault-tolerant RAID level (RAID 1, 3 or 5) is selected, the subsystem will start initializing parity.

Use the [ESC] key to view the status of the created logical drive.

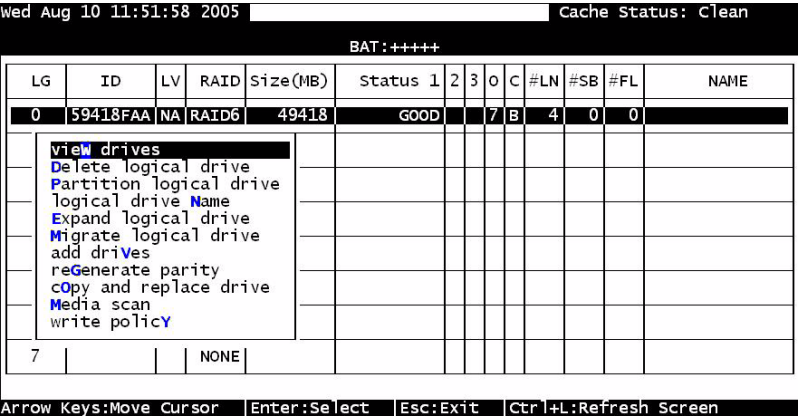


Fig. 7.15 Logical Drive Status Screen



NOTICE!
Only logical drives with RAID levels 1, 3 or 5 will take the time to initialize the logical drive. Logical drives with RAID level 0 and NRAID do not perform logical drive initialization; the drive initialization will be finished almost immediately.

7.5 Creating a Logical Volume

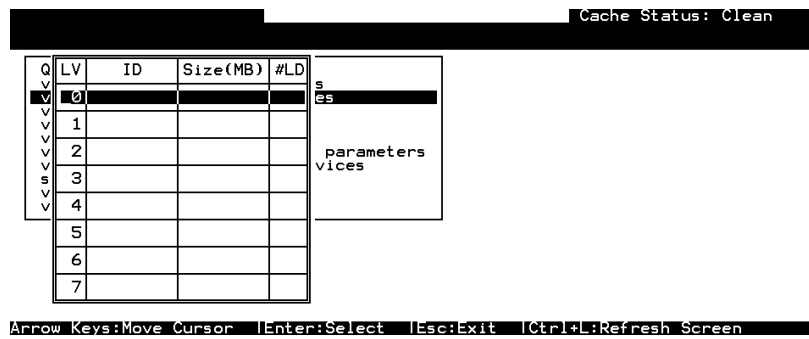


Fig. 7.16 View and Edit Logical Volumes Screen

A logical volume consists of one or several logical drives. Choose *View and Edit Logical Volumes* in the Main Menu to display the current logical volume configuration and status on the screen. Choose a logical volume number (0-7) that has not yet been defined, then press [ENTER] to proceed. A prompt “Create Logical Volume?” will appear. Select Yes and press [ENTER].

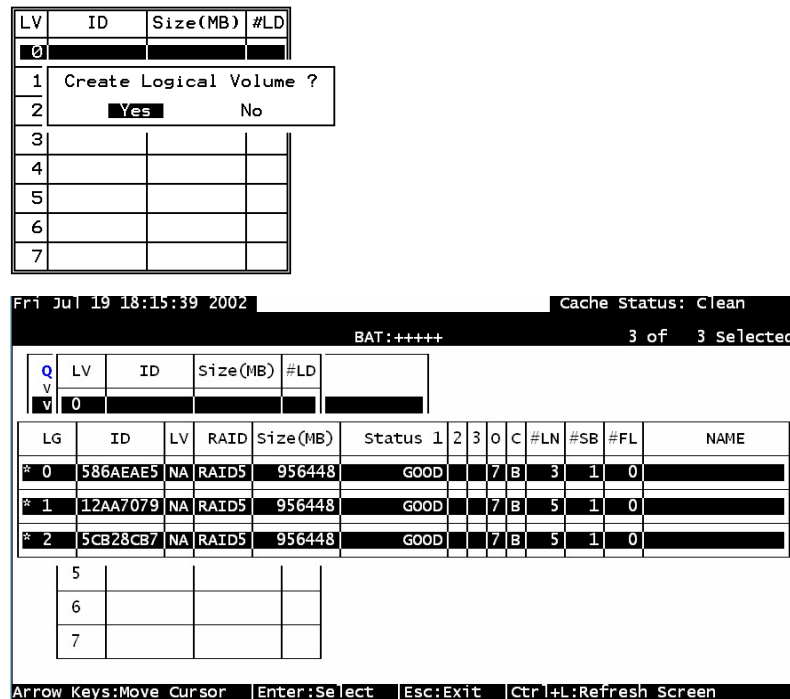


Fig. 7.17 Select Logical Drive Page

Select one or more logical drive(s) available on the list. The same as creating a logical drive, the logical drive(s) can be tagged for inclusion by positioning the cursor bar on the desired drive and pressing [ENTER] to select. An asterisk (*) will appear on the selected logical drive. Pressing [ENTER] again will deselect a logical drive.

Q	LV	ID	Size(MB)	#LD	
v	0				
v	1	Initialization Mode - On-Line			
v	2	Write Policy - Default			
v	3	Logical Volume Assignment - Primary			
s	4				
v	5				
v	6				
v	7				

Fig. 7.18 Selection Sub-menu

Use the arrow keys to select a sub-menu and change the initialization mode, write policy, or the managing controller.

Logical volumes can be assigned to different controllers (primary or secondary). The default is primary.

Logical Drive Count	: 1
Write Policy	: Default(Write-Back)
Create Logical Volume ?	
Yes	No

Fig. 7.19 Confirm Box

Note that if a logical volume is manually assigned to a specific controller, all its members' assignments will also be shifted to that controller.

When all the member logical drives have been selected, press [ESC] to continue. The confirm box displays. Choose Yes to create the logical volume.

Q	LV	ID	Size(MB)	#LD	
v	0	E7B465D	49418	2	
v	1	View logical drive			
v	2	Delete logical volume			
v	3	Partition logical volume			
v	4	Expand logical volume			
v	5	Write policy			
s	6				
v	7				

Fig. 7.20 View Logical Drive Selection Box

Press [ENTER] on a configured volume, and the information of the created logical volume displays.

LV:	Logical Volume ID
ID:	Unique ID for the logical volume, randomly generated by the RAID controller firmware
Size:	Capacity of this volume
#LD:	Number of the included members

7.6 Partitioning a Logical Drive/Logical Volume

The process of partitioning a logical drive is the same as that of partitioning a logical volume. The partitioning of a logical volume is used as an example in the proceeding discussion. Note that partitioning can be very useful when dealing with a very large capacity; however, partitioning a logical drive or logical volume is not a requirement for RAID configuration.

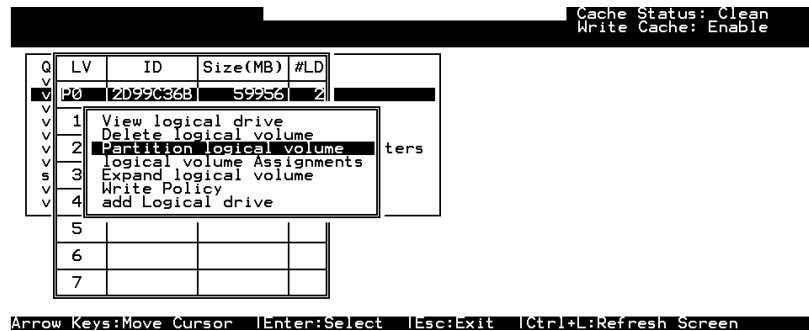


Fig. 7.21 Partition Logical Volume Page

Choose the logical volume you wish to partition, then press [ENTER]. Choose *Partition logical volume*, then press [ENTER]. Select from the list of undefined partitions and press [ENTER]. A list of partitions displays. If the logical volume has not yet been partitioned, all volume capacity will list as “partition 0.”

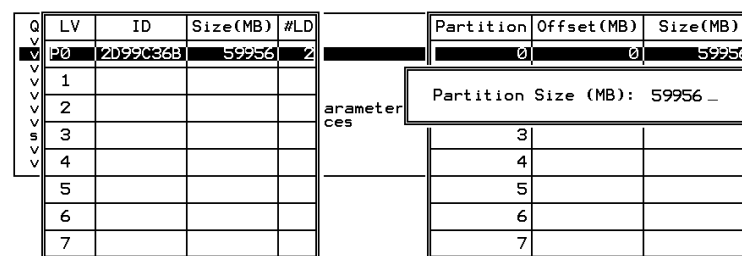


Fig. 7.22 Partition List

Press [ENTER] and type the desired size for the selected partition, and then press [ENTER] to proceed. The remaining size will be automatically allotted to the next partition. When prompted by the “Partition Logical Volume?” message, choose Yes to confirm then press [ENTER]. Follow the same procedure to partition the remaining capacity of your logical volume.




Fig. 7.23 Confirm Box

When a partition of a logical drive/logical volume is deleted, the capacity of the deleted partition will be added to the previous partition.

**WARNING!**

- Whenever a partition is changed, it is necessary to reconfigure all host LUN mappings. All data in it will be lost and all host LUN mappings will be removed when there is any change to the partition capacity.
- If operating in a Unix-based system, reset the subsystem for the configuration changes to take effect if any changes were made to partition sizes and partition arrangement.

7.7 Mapping a Logical Volume to Host LUNs

Select *View and Edit Host luns* in the Main Menu, then press [ENTER].

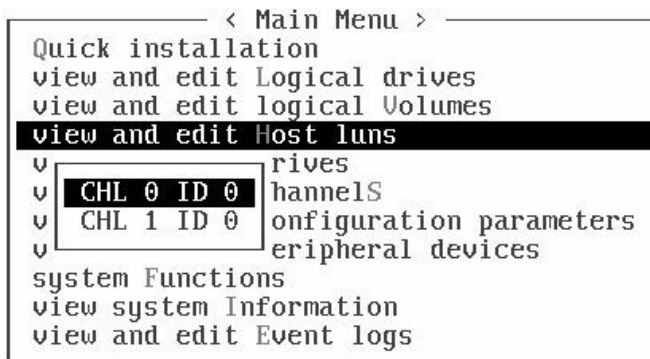


Fig. 7.24 View and Edit Host Luns Screen

A list of host channel/ID combinations appears on the screen. The diagram above shows two host channels and each is designated with a default ID. More can be added on each channel. Multiple IDs on host channels are necessary for redundant controller configuration. Details on creating multiple IDs and changing channel modes will be discussed later. Choose a host ID by pressing [ENTER].

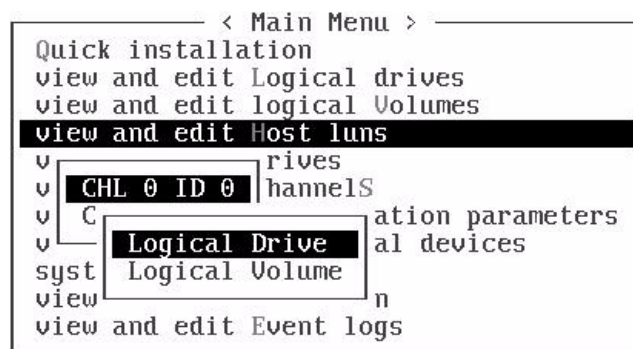


Fig. 7.25 Host Channel/ID List

Choose the channel-ID combination you wish to map, then press [ENTER] to proceed. Choose mapping a *Logical Drive* or a *Logical Volume* on the drop box.

Thu Sep 1 16:04:23 2005

Cache Status: Clean

BAT: NONE

< Main Menu >

Quick installation

view and edit Logical drives

view and edit Logical Volumes

view and edit IOSTatus

view CHL 0 ID 0 channels

view CHL 1 ID 0 configuration par

view peripheral device

System Functions

view system information

view and edit Event logs

LUN	LU/LD	DRV	Partition	Size(MB)	RAID
0					
1					
2					
3					
4					
5					
6					
7					

Arrow Keys: Move Cursor

Enter: Select

Esc: Exit

Ctrl-L: Refresh Screen

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0	6A62973B	NA	RAID	239112	GOOD						7	8	1	- 0

Partition	Offset(MB)	Size(MB)	3					
0	0	5000	4					
1	5000	10000	5					
2	15000	20000	6					
3	35000	204112	7					

Fig. 7.26 Mapping Result Page

7.8 Assigning a Spare Drive and Rebuild Settings

7.8.1 Adding Local Spare Drive

A spare drive is a standby drive automatically initiated by controller firmware to replace a failed drive. A spare drive must have an equal or larger capacity than the array members. A Local Spare should have a capacity equal to or larger than the members of the logical drive it is assigned to. A Global Spare should have a capacity equal to or larger than all physical drives in a RAID subsystem.

Quick view view view view view view view view view view	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	3	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
	4	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
	View drive information add Local spare drive add Global spare drive drive Utilities disk Reserved space - unformatted					tor 7Y250M0
						tor 7Y250M0
						tor 7Y250M0
	8	239367	150MB	NONE	NEW DRV	Maxtor 7Y250M0
	9	239367	150MB	1	ON-LINE	Maxtor 7Y250M0
	10	239367	150MB	1	ON-LINE	Maxtor 7Y250M0

Fig. 7.27 Add Local Spare Drive Page

Choose *View and Edit Drives* on the Main Menu, then press [ENTER]. Move the cursor bar to a drive that is not assigned to a logical drive or as a spare drive (usually indicated as a "New Drive"), and then press [ENTER].

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	586AEAE5	NA	RAID5	956448	GOOD				7	B	5	1	0	
P1	12AA7079	NA	RAID5	956448	GOOD				7	B	5	1	0	
P2	5CB28CB7	NA	RAID5	956448	GOOD				7	B	5	1	0	

Fig. 7.28 Logical Drives List

Choose *Add Local Spare Drive* and press [ENTER]. A list of logical drives displays. Move the cursor bar to a logical drive, then press [ENTER]. The unassigned disk drive will be associated with this logical drive as a Local Spare.

Add Local Spare Drive ?
Yes No

Fig. 7.29 Confirming Adding Local Spare Drive

When prompted to "Add Local Spare Drive?," choose Yes to confirm.

7.8.2 Adding a Global Spare Drive

A Global Spare replaces the failed drive in any logical drive of a RAID subsystem.

Quick view	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID	
view	3	239367	150MB	0	ON-LINE	Maxtor	7Y250M0
view	4	239367	150MB	0	ON-LINE	Maxtor	7Y250M0
view	View drive information					tor	7Y250M0
view	add Local spare drive					tor	7Y250M0
view	add Global spare drive					tor	7Y250M0
view	Add Global Spare Drive ?					tted	tor 7Y250M0
view	8	Yes No			RV	Maxtor	7Y250M0
	9	239367	150MB	1	ON-LINE	Maxtor	7Y250M0
	10	239367	150MB	1	ON-LINE	Maxtor	7Y250M0

Fig. 7.30 Add Global Spare Drive Page

Move the cursor bar to the drive that is not a member drive or a spare (usually indicated as a "New Drive"), and then press [ENTER]. Choose *Add Global Spare Drive*. When prompted to "Add Global Spare Drive?," choose Yes.

7.9 Viewing and Editing Logical Drives and Drive Members

Choose *View and Edit Logical Drives* in the Main Menu to display the array status. Refer to the previous chapter for more details on the legends used in the Logical Drive's status. To see the drive member information, choose the logical drive by pressing [ENTER].

Wed Sep 7 14:03:16 2005

Cache Status: Clean

BAT: NONE

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0	12E37A22	NA	RAID6	476428	GOOD					7	B	4	0	0
	view drives			428	GOOD					7	B	4	0	0
	Delete logical drive													
	Partition logical drive			642	GOOD					7	B	4	0	0
	logical drive Name													
	Expand logical drive													
	migrate logical drive													
	add drives													
	reGenerate parity													
	copy and replace drive													
	Media scan													
	write policy													
7			NONE											

Arrow Keys: Move Cursor | Enter: Select | Esc: Exit | Ctrl+L: Refresh Screen

Fig. 7.31 View and Edit Logical Drives Page

The disk drive-related functions include:

Selections	Function
View Drive	Display drive member information
Delete Logical Drive	To delete a logical drive
Partition Logical Drive	To delete a partition of a logical drive
Logical Drive Name	Assign a name to a logical drive
Rebuild Logical Drive	Rebuild a logical drive when a drive failed
Expand Logical Drive	To expand the size of a logical drive
Migrate Logical Drive	To migrate a logical drive to different RAID level
Add Drives	Add physical drive to a logical drive
Regenerate Parity	Regenerate logical drive parity
Copy and Replace Drive	Copy or replace a logical drive
Media Scan	Set media scan priority, iteration count and task schedules
Write Policy	Choose an appropriate write policy

7.9.1 Deleting a Logical Drive

Choose the logical drive you wish to delete, then press [ENTER]. Choose *Delete logical drive*, then choose Yes when prompted to confirm.

Wed Aug 10 13:49:24 2005Cache Status: Clean

BAT:++++

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME	
0	59418FAA	NA	RAID6	478734	GOOD						7	8	4	0	0

view drives

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	239367	133MB	0	ON-LINE	Maxtor 7Y250M0
2	239367	133MB	0	ON-LINE	Maxtor 7Y250M0
3	239367	133MB	0	ON-LINE	Seagate ST380023AS
4	239367	150MB	0	ON-LINE	WDC WD2500JD-00FYB0

7NONE

Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 7.32 Delete Logical Drive Page

7.9.2 Deleting a Partition of a Logical Drive

Q	LG	ID	LV	RAID	Size(MB)	Partition	Offset(MB)	Size(MB)	NAME
▼	P0	4149A729	NA	RAID5	39	0	0	3999	
▼	P1	76CD4DF6	NA	RAID0	119	1	3999	3999	
▼	2			NONE					
▼	3			NONE					
▼	4			NONE		4	15999	3999	
▼	5			NONE		5			
▼	6			NONE		6			
▼	7			NONE		7			

Partition Size (MB): 0

Fig. 7.33 Delete Partition Page

Choose the logical drive which has a partition you wish to delete, then press [ENTER]. Choose *Partition logical drive*. Partitions of the logical drive will be displayed in tabulated form. Move the cursor bar to the partition you wish to delete, then press [ENTER]. Enter “0” on the partition size to delete the partition.

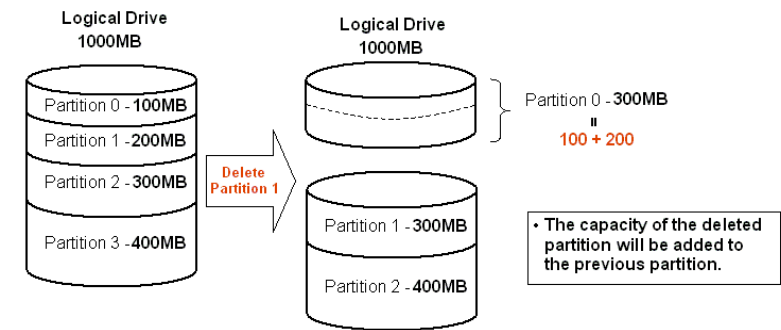


Fig. 7.34 Drive Space Allocated to the Previous Partition

As illustrated above, the capacity of the deleted partition will be added to the previous partition.



WARNING!

Whenever a partition is changed, it is necessary to reconfigure all host LUN mappings. All data kept in the partition and the host LUN mappings will be removed with any partition change.

7.9.3 Assigning a Name to a Logical Drive

Naming can help identify different arrays in a multi-array configuration. This function is also useful in special situations. For example, when one or more logical drives have been deleted, the array indexing is changed after system reboot. The second logical drive might become the first on the list.

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	477F9A88	0	RAID5	9999	GOOD				5		3	0	0	
P	View scsi drives			999	GOOD				5		2	-	0	
P	Delete logical drive			999	GOOD				5		2	-	0	
P	Partition logical drive			999	GOOD				5		2	-	0	
P	logical drive Name													
<div>Current Logical Drive Name: New Logical Drive Name: _</div>														
7			NONE											

Fig. 7.35 Name a Logical Drive Page

Choose the logical drive for which you wish to assign a name, then press [ENTER]. Choose *logical drive name*, then press [ENTER] again. The current name will be displayed. You may now enter a new name in this field. Enter a name, then press [ENTER] to save the configuration. The maximum number of characters for a logical drive name is 14.

7.9.4 Rebuilding a Logical Drive

If there is no spare drive in the system, a failed drive should be immediately replaced by a drive known to be good. Once the failed drive is replaced, the rebuild process can be manually initiated.

If you want the controller to auto-detect a replacement drive, make sure you have the following items set to “enabled”:


1. Periodic Drive Check Time
2. Periodic Auto-Detect Failure Drive Swap Check Time

These two configuration options can be found under “View and Edit Configuration Parameters” -> “Drive-Side Parameters.”

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
V	P0	4449A729	NA	RAID5	9998	DRU FAILED	R	2	0	0	
V	P	View scsi drives				GOOD	R	3	-	0	
V	P	Delete logical drive									
V	P	Partition logical drive									
V	P	logical drive Name									
V	P	Rebuild logical drive									
V	c	Rebuild Logical Drive ?									
S		Yes No									
6			NONE								
7			NONE								

Fig. 7.36 Rebuild Logical Drive Screen

Choose the logical drive that has a failed member drive, then press [ENTER]. Choose *Rebuild logical drive*, then press [ENTER]. When prompted to “Rebuild Logical Drive?,” select Yes.



The rebuild progress will be displayed.

Once rebuild has started, choose *Rebuild progress* to track its progress.

NOTICE!

The Rebuild function is only available when a logical drive with RAID level 1, 3 or 5 has a failed member. NRAID and RAID 0 configurations provide no data redundancy.



7.9.5 Expand Logical Drive

If there is an amount of unused capacity in a logical drive, the LD may be expanded. If there is no available unused capacity, then the LD cannot be expanded.

To expand a logical drive, first select *Expand Logical Drive* selection. Press [Enter] to perform the option.

7.9.6 Add Drive

More drives can be added into a logical drive to expand the logical drive capacity. To add a drive, first choose *Add Drive* selection and press [Enter]. All available disks will be shown on the screen. Use the Up and Down arrow keys to select the drive you wish to add then press [Enter]. When the drive is selected, a start “*” sign will appear in front of the slot number. Press [ESC], and then click Yes to confirm the action.

Wed Sep 7 15:13:06 2005										Cache Status: Clean															
BAT: NONE										0 of 4 Selected															
LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME											
0	12E37A22	NA	RAID6	476428	DRV FAILED					7 B	3	1	1												
1	281DFB99	NA	RAID6	476428	GOOD					7 B	4	1	0												
view drives				642	GOOD					7 B	4	1	0												
Delete Part Logi Expa migr add rege copy Medi writ	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID																			
	3	238469	150MB	NONE	USED DRV	WDC WD2500SD-01KCB0																			
	14	238469	150MB	NONE	NEW DRV	WDC WD2500SD-01KCB0																			
	15	238469	150MB	NONE	NEW DRV	WDC WD2500SD-01KCC0																			
	16	238469	150MB	NONE	NEW DRV	WDC WD2500SD-01KCB0																			

Fig. 7.37 Add Drive Screen

7.9.7 Regenerating Logical Drive Parity

(Applies to RAID Levels 1, 3 or 5)

If no verifying method is applied to data writes, this function can often be performed to verify parity blocks of a selected array. This function compares and recalculates parity data to correct parity errors.

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	4295529B	NA	RAID5	476000	GOOD					4 B	3	0	0	
Execute Regenerate Logical Drive Parity														
Overwrite Inconsistent Parity - Enabled														
Generate Check Parity Error Event - Enabled														
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

Fig. 7.38 Confirm Box

Choose the logical drive that you want to regenerate the parity for, and then press [ENTER]. Choose *Regenerate Parity*, then press [ENTER]. When prompted to “Regenerate Parity?,” select Yes.

Please refer to *Section 10 Data Integrity* for more information on parity regeneration.

7.9.8

Media Scan

Media Scan examines drives and detects the presence of bad blocks. If any data blocks have not been properly committed and are found during the scanning process, data from those blocks are automatically recalculated, retrieved and stored onto undamaged sectors. If bad blocks are encountered on yet another drive during the rebuild process, the block LBA (Logical Block Address) of those bad blocks will be shown. If rebuild is carried out under this situation, rebuild will continue with the unaffected sectors, salvaging the majority of the stored data.

Thu Jun 16 18:31:36 2005										Cache Status: Clean									
BAT:+++++																			
LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME					
P0	6E448ADE	NA	NRAID	238214	GOOD				7	0	1	-	0						
Media Scan Priority - Normal																			
Iteration Count - Single Time																			
Task Scheduler																			
3			NONE																
4			NONE																
5			NONE																
6			NONE																
7			NONE																

Arrow Keys:Move Cursor										!Enter:Select					!Esc:Exit					!Ctrl+L:Refresh Screen				
------------------------	--	--	--	--	--	--	--	--	--	---------------	--	--	--	--	-----------	--	--	--	--	------------------------	--	--	--	--

Arrow Keys:Move Cursor iEnter:Select iEsc:Exit iCtrl+L:Refresh Screen

Fig. 7.39 Media Scan Screen

There are two options with performing the Media Scan:

1. Media Scan Priority: determines how much of the system resources will be used for the drive scanning and recalculating process.
2. Iteration Count: determines how many times the scan is performed. If set to “continuous,” the scan will run in the background continuously until it is stopped by a user.

Media Scan Priority - Normal
Iteration Count - Single Time
Set Media Scan Iteration Count to Continuous Scan ?
Yes No

The system can automatically perform a Media Scan according to a preset task schedule. For more details, please refer to *Section 10 Data Integrity*.

7.9.9

Write Policy

view drives	
Delete logical drive	
Partition logical drive	
logical drive Name	
Expand logical drive	
M	
a	
r	
c	
M	
Change Write Policy ?	
Default Write-Back Write-Through	
write policy	

Fig. 7.40 Write Policy Screen

The Write-back cache setting can be configured differently on each array. Setting to the default value means the array setting is coordinated with the controller's general setting. The controller's general setting option can be found in “View and Edit Configuration Parameters” -> “Caching Parameters” -> “Write-Back Cache.” Note that cached writes are lost if a power failure occurs.

7.10 Viewing and Editing Host LUNs

Viewing or Deleting LUN Mappings

Choose the host channel and host ID combination you wish to view or delete.

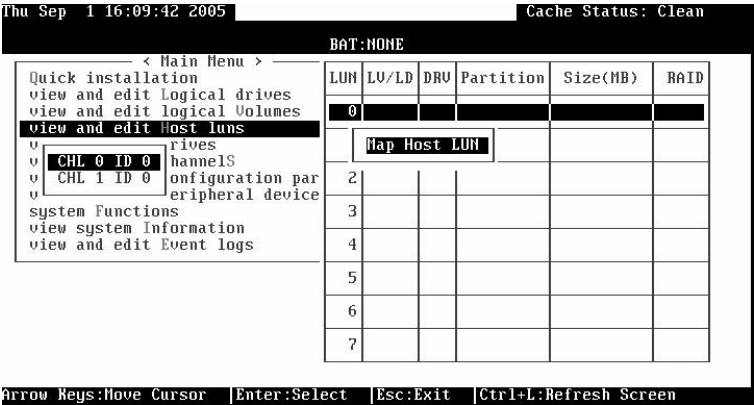


Fig. 7.41 View and Edit Host Luns Page

A list of the current LUN mapping will be displayed on the screen. Move the cursor bar to the LUN mapping you wish to delete, then press [ENTER]. Select Yes to delete the LUN mapping, or No to cancel.

7.11 Viewing and Editing Drives

Sat Jul 20 12:42:35 2002 Cache Status: Clean

BAT:++++

Quick view view view view view syst view view	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	1	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
	2	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
	3	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
	4	239367	150MB	GLOBAL	STAND-BY	Maxtor 7Y250M0
	5	239367	150MB	NONE	FRMT DRV	Maxtor 7Y250M0
	6	239367	150MB	1	ON-LINE	Maxtor 7Y250M0
	7	239367	150MB	1	ON-LINE	Maxtor 7Y250M0
	8	239367	150MB	NONE	NEW DRV	Maxtor 7Y250M0

Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 7.42 View and Edit Drives Main Menu

Choose *View and Edit Drives* in the Main Menu. All drives attached to the drive channels will be displayed on the screen.

Press [ENTER] on a selected drive to view a sub-menu featuring functions related to the individual hard drive.

Thu Sep 8 08:31:10 2005 Cache Status: Clean

BAT:NONE

Quick view view view view view syst view view	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
	1	238469	150MB	1	ON-LINE	WDC WD2500SD-01KCB0
		View drive information				WDC WD2500SD-01KCB0
		Identify drive				WDC WD2500SD-01KCB0
		clone Failing drive				WDC WD2500SD-01KCB0
		disk Reserved space - 256 mb				WDC WD2500SD-01KCB0
		media error reAssign count - 0				WDC WD2500SD-01KCB0
	5	238469	150MB	1	ON-LINE	WDC WD2500SD-01KCB0
	6	238469	150MB	1	ON-LINE	WDC WD2500SD-01KCB0
	7	238469	150MB	NONE	FRMT DRV	WDC WD2500SD-01KCC0
	8	238469	150MB	NONE	FRMT DRV	WDC WD2500SD-01KCB0

Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 7.43 View Drive Information Sub-menu

The disk drive-related functions include:

Function	Availability
View Drive Information	Every drive type
Identify Drives	Every drive type
Clone Failing Drive	Members of logical drives
Delete Global/Local Spare Drive	Global/Local Spare drive
Add Global/Local Spare Drive	Non-configured drives
Media Scan	Non-configured drives; Media Scan on drive members can be found in View and Edit Logical Drives
Drive Utilities: Read/Write Test	New drives
Disk Reserved Space	View only on member drives, reserved space removable on the used or formatted drives
Media Error Reassign Count	View the number of media error encountered on each drive

Using Media see *Section 7.9.8 Media Scan*. Details on the use of Clone Failing Drive can be see *Section 13 Implementations for AV Applications*.

7.11.1 Deleting Spare Drive (Global / Local Spare Drive)

Move the cursor to a Local Spare Drive or Global Spare Drive, then press [ENTER]. Choose *Delete Global/Local Spare Drive*, then press [ENTER] again. Choose Yes to confirm.

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID	
1	239367	150MB	0	ON-LINE	Maxtor	7Y250M0
2	239367	150MB	0	ON-LINE	Maxtor	7Y250M0
3	239367	150MB	0	ON-LINE	Maxtor	7Y250M0
4	239367	150MB	GLOBAL	STAND-BY	Maxtor	7Y250M0
					View drive information	
					Delete global/local spare drive	
					Delete Spare Drive ?	
					Yes No	
8				EW DRV	Maxtor	7Y250M0

Fig. 7.44 Delete Global/Local Spare Drive Page



NOTICE!

The spare drive you deleted or any drive you replaced from a logical unit will be indicated as a "used drive."

7.11.2 Disk Reserved Space

The 256MB of reserved space can be removed from a drive once the drive is excluded from a logical drive. The reserved space, a space formatted with a micro-file system, can also be manually created from a new disk drive.

5	239367	133MB	NONE	FRMT DRV	Maxtor	7Y250M0
					View drive information	
					add Local spare drive	
					add Global spare drive	
					disk Reserved space - 256 mb	
					Remove Disk Reserved Space ?	
					Yes No	

Fig. 7.45 Remove Disk Reserved Space Screen

View drive information	
add Local spare drive	
add Global spare drive	
drive Utilities	
disk Reserved space - unformatted	
256MB	
Format Disk Reserved Space ?	
Yes	No

Fig. 7.46 Format Disk Reserved Space Screen

7.11.3 Identifying Drives

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	239367	133MB	0	ON-LINE	Maxtor 7Y250M0
View drive information add Local spare drive add Global spare drive Identify drive					tor 7Y250M0
flash All drives flash Selected drive flash all But selected drive					tor 7Y250M0
6					DC WD740GD-00FLX0
7				ABSENT	
8				ABSENT	

Fig. 7.47 Identify Drive Page

Whenever there is a failed drive in a logical drive, make it a point to replace the failed drive with a new, healthy drive to keep the logical drive working.

If you replaced the wrong drive when trying to replace a failed drive you will no longer be able to access the logical drive because you have inadvertently failed another drive.

To prevent this from happening, the controller provides an easy way to identify the faulty drive. By forcing certain drive LEDs to light for a configurable period of time, the faulty drive can be identified, and thus reduce the chance of removing the wrong drive. This function can be especially helpful in an installation site operating with hundreds of drives.

Flash Selected Drive

The Read/Write LED of the drive you selected will light steadily for a configurable period of time, from 1 to 999 seconds.




Fig. 7.48 Flash Selected Drive

Flash All Drives

The Read/Write LEDs of all connected drives will light for a configurable period of time. If the LED of the defective drive did not light on the “Flash Selected Drive” function, use “Flash All Drives” to verify the fault. If the “Flash All Drives” function is executed, and the defective drive’s LED still does not respond, it can be a drive tray problem or the drive is dead.




Fig. 7.49 Flash All Drives

Flash All but Selected Drives

Except the selected drive, tExEhe Read/Write LEDs of all connected drives will light for a configurable period of time ranging from 1 to 999 seconds. If an administrator cannot be sure of the exact location of a specific drive, this function will help to indicate where it is. This can prevent removal of the wrong drive when a drive fails and is about to be replaced.




Fig. 7.50 Flash All but Selected Drives

The drive-identifying function can be selected from “Main Menu”/”View and Edit Drives”/”Identify Drives.”

7.11.4

Media Error Reassign Count

Thu Sep 8 08:48:10 2005

Cache Status: Clean

BAT: NONE

Quick view view view view view view view view view	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID	
	1	238469	150MB	0	ON-LINE	WDC	WD2500SD-01KCB0
	View drive information					WDC	WD2500SD-01KCB0
	Identify drive					WDC	WD2500SD-01KCB0
	clone Failing drive					WDC	WD2500SD-01KCB0
	disk Reserved space - 256 mb					WDC	WD2500SD-01KCB0
	media error reassign count - 0					WDC	WD2500SD-01KCB0
	5	238469	150MB	NONE	FRMT DRV	WDC	WD2500SD-01KCB0
	6	238469	150MB	NONE	FRMT DRV	WDC	WD2500SD-01KCB0
	7	238469	150MB	NONE	FRMT DRV	WDC	WD2500SD-01KCC0
	8	238469	150MB	NONE	FRMT DRV	WDC	WD2500SD-01KCB0

Arrow Keys: Move Cursor | Enter: Select | Esc: Exit | Ctrl+L: Refresh Screen

Arrow Keys: Move Cursor |Enter: Select |Esc: Exit |Ctrl+L: Refresh Screen

Fig. 7.51 Media Error Reassign Count Page

This option provides information about the occurrence of drive's media error. If a disk drive has encountered media errors, users may enable drive S.M.A.R.T. function to detect drive failure status. If messages show media error come frequent, the target disk drive might have reached its life expectancy, and should be replaced immediately.

7.12 Viewing and Editing Channels

Thu Sep 1 16:23:13 2005Cache Status: Clean

BAT:NONE

Q	Ch1	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
U	0	Host	0	160.0MHz	Wide	L	On	Async	Narrow
U	1	Host	0	160.0MHz	Wide	L	On	Async	Narrow
U	2	Drive		AUTO	SATA				SATA
U	3	Drive		AUTO	SATA				SATA
S	4	Drive		AUTO	SATA				SATA
U	5	Drive		AUTO	SATA				SATA
U	6	Drive		AUTO	SATA				SATA
U	7	Drive		AUTO	SATA			1.33GHZ	SATA

Arrow Keys:Move CursorEnter:SelectEsc:ExitCtrl+L:Refresh Screen

Fig. 7.52 View and Edit Channels Page

Choose *View and Edit Channels* in the Main Menu to display channel status.

Thu Sep 1 16:21:38 2005Cache Status: Clean

BAT:NONE

Q	Ch1	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
U	0	Host	0	160.0MHz	Wide	L	On	Async	Narrow
U	1	channel Mode			Wide	L	On	Async	Narrow
U	2	view and edit scsi id			SATA				SATA
U	3	scsi terminator			SATA				SATA
U	4	sync transfer Clock			SATA				SATA
S	5	Wide transfer			SATA				SATA
U	6	parity check - Enabled							
U	7	view chip information						1.33GHZ	SATA

Arrow Keys:Move CursorEnter:SelectEsc:ExitCtrl+L:Refresh Screen

Fig. 7.53 View and Edit ID Page

7.12.1 Viewing and Editing IDs - Host Channel

0	Host	0	160.0MHz	Wide	L	On	Async	Narrow	ID 1
1	ID 0	0	160.0MHz	Wide	L	On	Async	Narrow	ID 2
2	Add Channel SCSI ID			ATA				SATA	ID 3
3	Delete Channel SCSI ID			ATA				SATA	ID 4
4	Drive		AUTO	SATA				SATA	ID 5
5	Drive		AUTO	SATA				SATA	ID 6
6	Drive		AUTO	SATA				SATA	ID 7
7	Drive		AUTO	SATA			1.33GHZ	SATA	ID 8

ID 9
ID 10
ID 11
ID 12
ID 13
ID 14
ID 15

Fig. 7.54 Existing ID List

Choose a host channel, then press [ENTER]. Choose *View and Edit ID*. A list of existing ID(s) will be displayed on the screen. Select one of the existing IDs and press [ENTER]. You may then choose to add or delete an existing ID.

7.12.2

Adding an ID

0	Host	0	160.0MHz	Wide	L	On	Async	Narrow	ID 1
1	ID 0	t	0	160.0MHz	Wide	L	On	Async	ID 2
2	Add Channel SCSI ID			ATA				SATA	ID 3
3	Delete Channel SCSI ID			ATA				SATA	ID 4
4	Drive		AUTO	SATA				SATA	ID 5
5	Drive		AUTO	SATA				SATA	ID 6
6	Drive		AUTO	SATA				SATA	ID 7
7	Drive		AUTO	SATA			1.33GHZ	SATA	ID 8
									ID 9
									ID 10
									ID 11
									ID 12
									ID 13
									ID 14
									ID 15

Fig. 7.55 Add ID Screen

Press [ENTER] on one of the existing IDs. Choose *Add Channel ID*, and then choose to assign an ID. A list of host IDs will appear. Choose an ID. DO NOT choose an ID used by another device on the same channel.

7.12.3

Deleting an ID

Thu Sep 1 16:24:59 2005									
Cache Status: Clean									
BAT:NONE									
Q	Ch1	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
U	0	Host	0	160.0MHz	Wide	L	On	Async	Narrow
U	1	ID 0	t	0	160.0MHz	Wide	L	On	Async
U	2	Add Channel SCSI ID			ATA				SATA
U	3	Delete Channel SCSI ID			ATA				SATA
S	4	Delete SCSI ID 0 ?			SATA				SATA
U	5	Yes No			SATA				SATA
U	6	Drive		AUTO	SATA				SATA
U	7	Drive		AUTO	SATA			1.33GHZ	SATA

Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 7.56 Delete ID Screen

Choose the host bus ID you wish to delete. Choose *Delete Channel ID*. The dialog box "Delete ID#?" will appear. Select Yes, then press [ENTER] to confirm.

NOTICE!

- Every time you change a channel ID, you must reset the subsystem/controller for the changes to take effect.
- Multiple target IDs can be applied to a host channels while each drive channel has only one ID.
- At least one controller's ID has to be present on each channel bus.



7.12.4 Data Rate (Channel Bus)

Thu Sep 1 16:18:06 2005

Cache Status: Clean

BAT: NONE

Q	Ch1	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid	
V	0	Host	0	160.0MHz	Wide	L	On	Async	Narrow	160.0MHz
V										80.0MHz
V										40.0MHz
V										33.0MHz
V	1	channel Mode			Wide	L	On	Async	Narrow	20.0MHz
V	2	view and edit scsi Id			SATA				SATA	16.6MHz
V	3	scsi Terminator			SATA				SATA	13.8MHz
V	4	sync transfer Clock			SATA				SATA	10.0MHz
S		Wide transfer								8.0MHz
V		parity check - Enabled								6.7MHz
V		view chip inFormation								5.8MHz
	5	Drive		AUTO	SATA				SATA	5.0MHz
	6	Drive		AUTO	SATA				SATA	4.0MHz
	7	Drive		AUTO	SATA			1.33GHZ	SATA	3.3MHz
										2.8MHz
										2.5MHz
										Async

Arrow Keys: Move Cursor | Enter: Select | Esc: Exit | Ctrl+L: Refresh Screen

Fig. 7.57 Data Rate Screen

This option is available in the configuration menu of Fibre host channel and the drive channel configuration menus of Fibre- or SATA-based subsystems. Default is “AUTO” and should work fine with most disk drives. Changing this setting is not recommended unless some particular bus signal issues occur.

Most SATA/ATA-based systems connect only one drive per SATA/ATA channel (4 for multi-lane with SATA II) to help avoid a single drive failure from affecting other drives. The maximum mechanical performance of today’s drives can reach around 30MBps (sustained read) which is still far below the bandwidth of a drive channel bus. Setting the SATA bus speed to a lower value can get around some problems, but will not become a bottleneck to system performance.

Note that the SATA/ATA speed is the maximum transfer rate of the SATA/ATA bus in that mode. It does not mean the drive can actually carry out that amount of sustained read/write performance. For the performance of each drive model, please refer to the documentation provided by drive manufacturer.

7.12.5 View Chip Information

Q	Ch1	Mode	ID	DefSynClk	DefWid	S	Term	CurSynClk	Curwid
V	0	Host	112	AUTO	Serial	F	NA		
V	1	view and edit Id					NA		
V	2	view chip inFormation						1.33GHZ	SATA
V	3	Chip Type			Tachyon DX4	n)		1.33GHZ	SATA
S		Chip Rev. ID			5				
V		Chip Fw Rev. ID			N/A			1.33GHZ	SATA
	4								
	5	Drive		AUTO	SATA			1.5GHZ	SATA
	6	Drive		AUTO	SATA			1.33GHZ	SATA
	7	Drive		AUTO	SATA			1.5GHZ	SATA

Fig. 7.58 View Chip Information

This is a view only option showing basic information about the SCSI-320 Channel chip processor.

7.13 System Functions

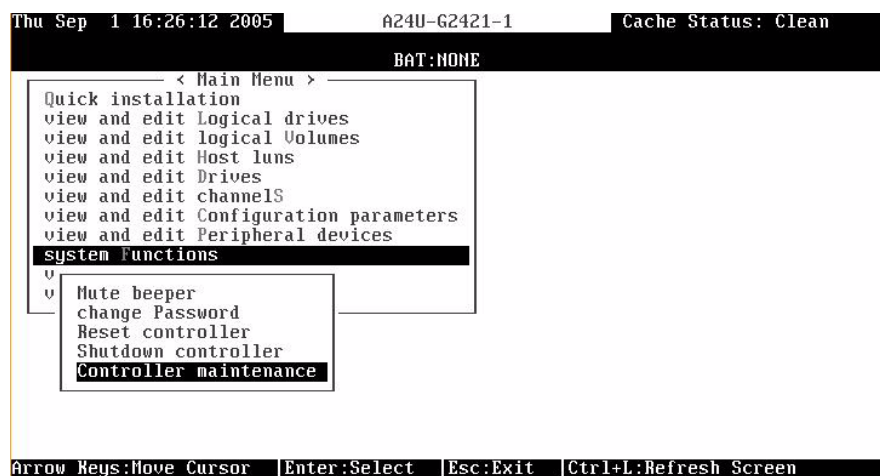


Fig. 7.59 System Functions Page

Choose *System Functions* in the Main Menu, then press [ENTER] to display the System Functions menu. Move the cursor bar to an item, then press [ENTER].

7.13.1 Mute Beeper

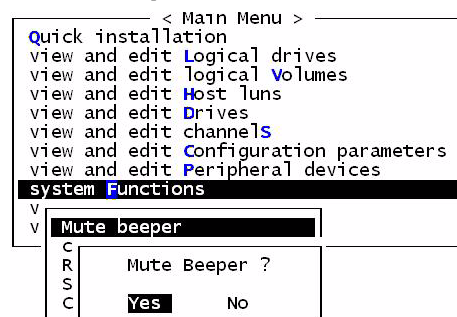


Fig. 7.60 Mute Beeper Screen

When the subsystem's beeper (onboard alarm) has been activated, choose *Mute Beeper*," then press [ENTER]. Choose Yes and press [ENTER] in the next dialog box to turn the beeper off temporarily for the current event. The beeper will still be activated by the next event.

7.13.2 Change Password

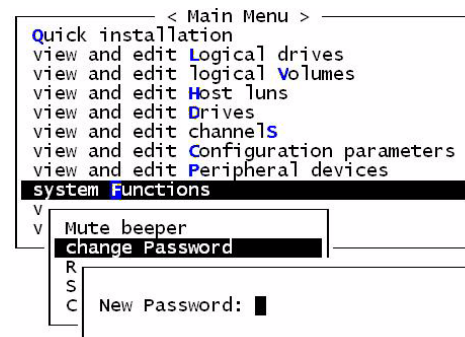


Fig. 7.61 Change Password Screen

Use the subsystem's password to protect the array from unauthorized entry. Once the subsystem password has been set, regardless of whether the front panel, the RS-232C terminal

interface or RAIDWatch Manager is used, you can only access the subsystem by providing the correct password.

NOTICE!

- The controller verifies the password when entering the Main Menu from the initial screen or when making a configuration change. If the controller is going to be left unattended, the “Password Validation Timeout” can be set to “Always Check.” Setting the validation timeout to “Always Check” will protect the controller configuration from any unauthorized access.
- The controller password and controller name share a 16-character space. The maximum numbers of characters for the controller password is 32. If the controller name occupies 32 characters, there is only one character left for the controller password, and vice versa.

7.13.3

Changing the Password

To set or change the controller password, move the cursor bar to *Change Password*, then press [ENTER].

If a password has previously been set, the controller will ask for the old password first. If the password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

Key-in the old password, then press [ENTER]. If the password is incorrect, it will not allow you to change the password. Instead, it will display the message “Password incorrect!,” then go back to the previous menu.

If the password is correct, or there is no preset password, it will ask for the new password.

7.13.4

Setting a New Password

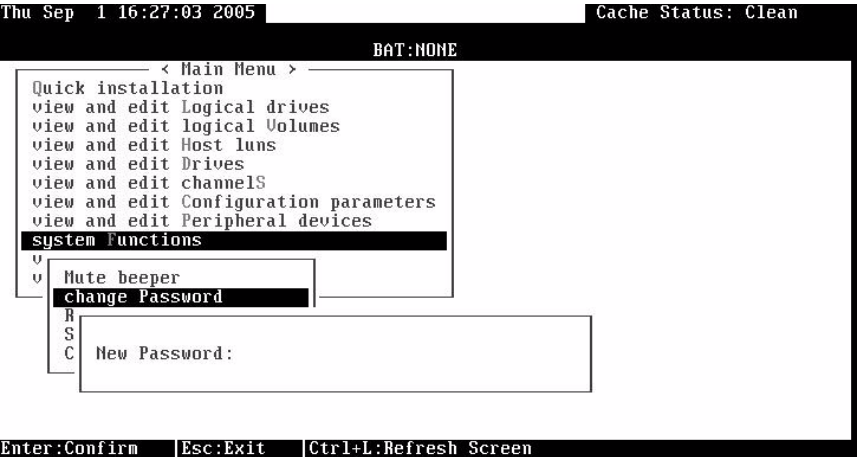


Fig. 7.62 Set New Password Screen

Enter the desired password in the column, then press [ENTER]. The next dialog box will display “Re-Enter Password.” Enter the password again to confirm and press [ENTER].

The new password will now become the controller’s password. Providing the correct password is necessary when entering the Main Menu from the initial screen.

7.13.5

Disabling the Password

To disable or delete the password, press [ENTER] in the empty column that is used for entering a new password. The existing password will be deleted. No password checking will occur when entering the Main Menu or when making a configuration change.

7.13.6

Reset Controller

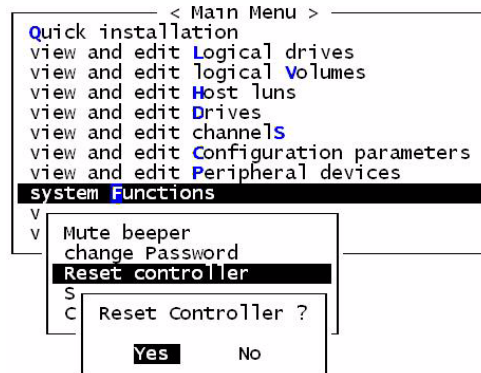


Fig. 7.63 Reset Controller Screen

To reset the controller without powering off the system, move the cursor bar to *Reset Controller*, then press [ENTER]. Choose *Yes* in the dialog box that follows, then press [ENTER]. The controller will now reset as well as power off or re-power on.

7.13.7

Shutdown Controller

Before powering off the controller, unwritten data may still reside in cache memory. Use the “Shutdown Controller” function to flush the cache content. Move the cursor bar to *Shutdown Controller*, then press [ENTER]. Choose *Yes* in the dialog box that follows, then press [ENTER]. The controller will now flush the cache memory.

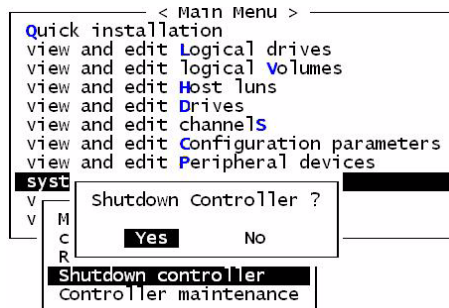


Fig. 7.64 Shutdown Controller Screen

For Controller Maintenance functions, such as Download Firmware, see *Section 15 System Functions - Upgrading Firmware*.

7.13.8

Adjust LCD Contrast

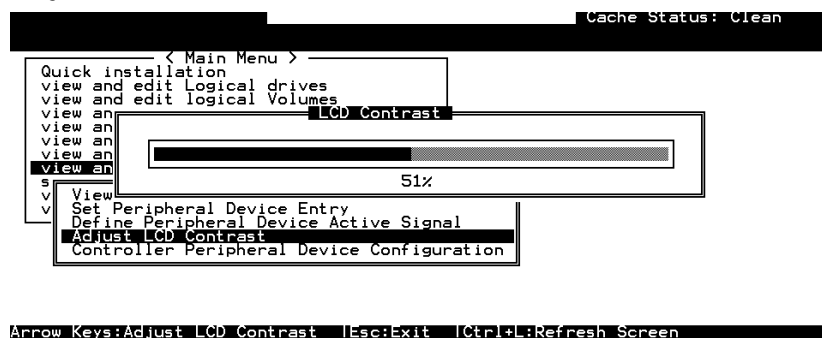


Fig. 7.65 Adjust LCD Contrast Page

The controller LCD contrast is set at the factory to a level that should be generally acceptable. If changes are required, the controller is equipped with an LCD contrast adjustment circuit,

allowing the contrast to be adjusted either via the RS-232 Terminal Emulation Menus or by the LCD User Interface.

7.14 Controller Parameters

7.14.1 Controller Name

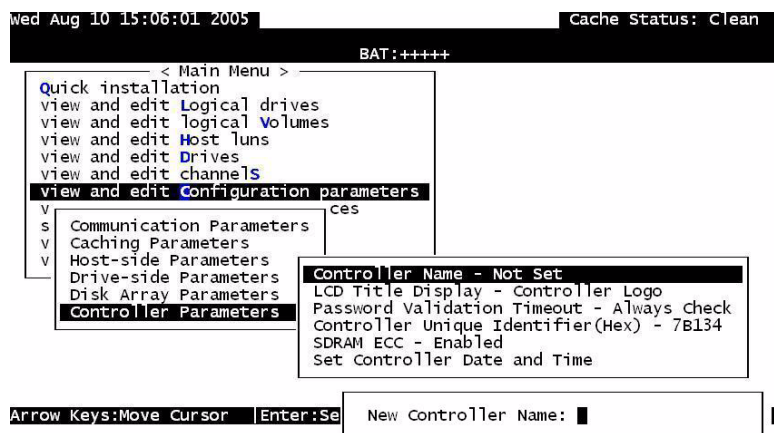


Fig. 7.66 View and Edit Configuration Parameters Page

Choose *View and Edit Configuration Parameters, Controller Parameters*, then press [ENTER].

The current name displays. Press [ENTER]. Enter a name in the dialog box that prompts, then press [ENTER].

7.14.2 LCD Title Display - Controller Name

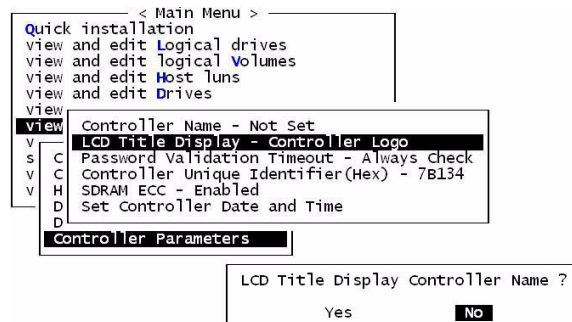


Fig. 7.67 Controller Parameters Screen

Choose *View and Edit Configuration Parameters, Controller Parameters*, then press [ENTER].

Choose to display the embedded controller logo or any given name on the LCD. Giving a specific name to each controller will make them easier to identify if you have multiple RAID systems remotely monitored.

7.14.3 Saving NVRAM to Disks

You can choose to backup your controller-dependent configuration information to disks. We recommend using this function to save configuration information whenever a configuration change is made. The information will be duplicated and distributed to all logical configurations of drives.

At least a RAID configuration must exist for the controller to write your configuration data onto it.

From the Main Menu, choose *System Functions*. Use the arrow keys to scroll down and select *Controller Maintenance*, *Save NVRAM to Disks*, then press [ENTER].

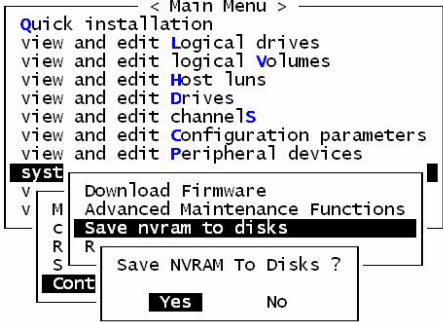


Fig. 7.68 Save NVRAM to Disk Screen

Choose Yes to confirm.

A prompt will inform you that NVRAM information was successfully saved.

7.14.4 Restore NVRAM from Disks

Use this function to restore your NVRAM information that you previously saved onto disk.

From the Main Menu, choose *System Functions*. Use the arrow keys to scroll down and select *Controller Maintenance*, *Restore NVRAM from disks*, and then press [ENTER].




Fig. 7.69 Restore NVRAM from Disks Screen

Press [ENTER] on one of the selections to confirm.

A prompt will notify you that the controller NVRAM data was successfully restored from disks.

7.14.5 Password Validation Timeout

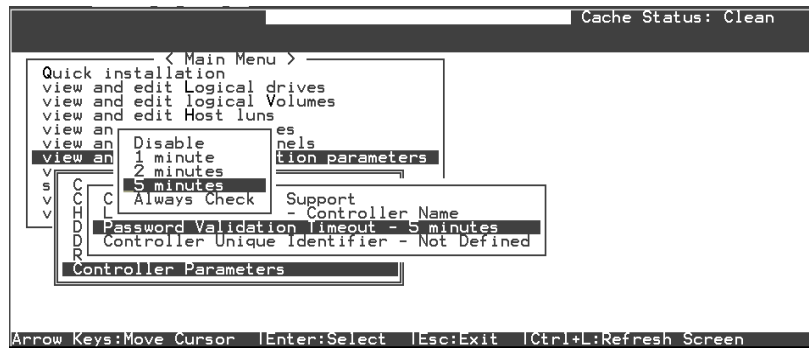


Fig. 7.70 Password Validation Timeout Page

Choose *View and Edit Configuration Parameters, Controller Parameters*, then press [ENTER]. Select *Password Validation Timeout*, and press [ENTER]. Choose to enable a validation timeout from “1 minute” to “Always Check.” The Always Check timeout will disable any configuration change made without entering the correct password.

7.14.6 Controller Unique Identifier

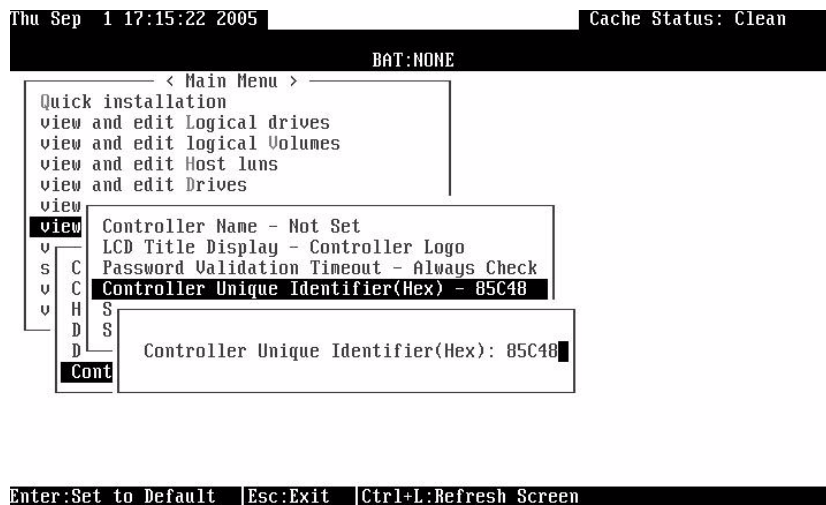


Fig. 7.71 View and Edit Configuration Parameters

Enter any hex number between “0” and “FFFFF” for the unique identifier. The value you enter *MUST* be different for each controller.

Enter a unique ID for every RAID subsystem whether it is configured in a single-controller configuration. The unique ID is necessary for the following:

1. MAC addresses for the controller's Ethernet port that will be taken over by a surviving controller in the event of single RAID controller failure.
2. The unique identifier setting can be accessed from "View and Edit Configuration Parameters" -> "Controller Parameters" -> "Controller Unique ID."

7.14.7 Set Controller Date and Time

This sub-menu only appears when the controller is equipped with a real-time clock.

Time Zone

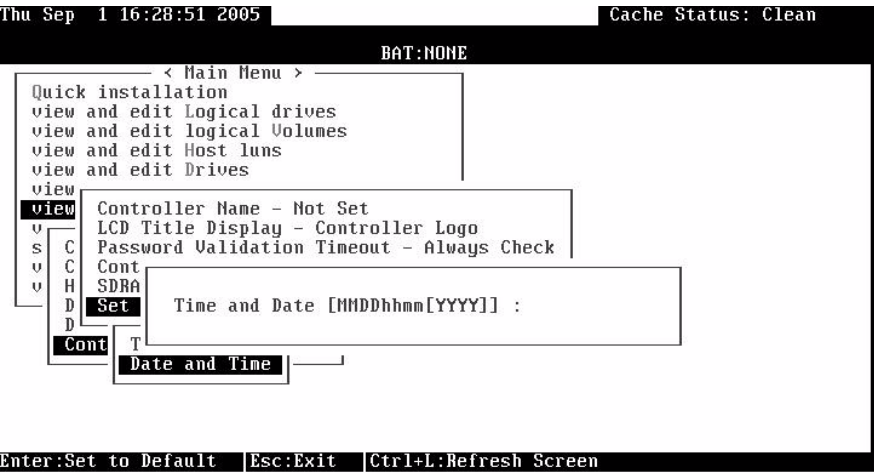


Fig. 7.72 Set Controller Date and Time Page

The controller uses GMT (Greenwich Mean Time), a 24-hour clock. To change the clock to your local time zone, enter the numbers of hours later than the Greenwich Mean Time after a plus (+) sign. For example, enter “+9” for Japan’s time zone.

Date and Time




Fig. 7.73 Enter Date and Time Screen

Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

7.15 View Drive Information

sat Jul 20 17:26:58 2002 Cache Status: Clean

BAT:++++

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
2	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
3	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
4	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
5	239367	150MB	NONE	NEW DRV	Maxtor 7Y250M0
6	239367	150MB	1	ON-LINE	Maxtor 7Y250M0
7	239367	150MB	1	ON-LINE	Maxtor 7Y250M0
8	239367	150MB	NONE	NEW DRV	Maxtor 7Y250M0

Esc:Exit | Ctrl+L:Refresh Screen

Fig. 7.74 View Drive Information Page

From the “View and Edit Drives” menu, select the drive that the utility is to be performed on, then press [ENTER]. Select *View drive information*, then press [ENTER].

7.16 Drive Utilities

7.16.1 Drive Read/Write Test

From the “View and Edit Drives” menu, select a new or used drive that the utility is to be performed on, then press [ENTER]. Select *Drive Utilities*, then press [ENTER]. Choose *Read/Write Test* and press [ENTER]. You can choose to enable/disable the following options:

1. Abort When Error Occurs
2. Drive Test for - Read Only/Read and Write
3. Execute Drive Testing

wed Aug 10 15:16:31 2005 Cache Status: Clean

BAT:++++

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
2	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
3	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
4	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
5	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
6	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
7	239367	150MB	0	ON-LINE	Maxtor 7Y250M0
8	239367	150MB	0	ON-LINE	Maxtor 7Y250M0

Esc:Exit | Ctrl+L:Refresh Screen

Fig. 7.75 Drive Read/Write Test Page

When finished with the configuration, select *Execute Drive Testing* and press [ENTER] to proceed.

The Read/Write test progress will be indicated by a status bar.




Fig. 7.76 Disk Read/Write Testing Completion Percentage

You may press [ESC] and select "Read/Write Test" later and choose to "View Read/Write Testing Progress." If you want to stop testing the drive, select *Abort Drive Testing* and press [ENTER] to proceed.




Fig. 7.77 View Read/Write Testing Progress Screen



NOTICE!
The disk drive on which a read/write test will be performed cannot be a spare drive (local or global) or a member of a logical drive. The "Read/Write Test" option will not appear if the drive is not indicated as a "New Drive" or a "Used Drive." Also, a drive formatted with a 256MB reserved space is also excluded from selection.

8 Host-side and Drive-side Parameters

This chapter discusses the advanced options for configuring and maintaining a RAID system. Each function is given a brief explanation as well as a configuration sample. Terminal screens are used in the configuration samples. Some of the operations require basic knowledge of RAID technology and are only recommended for an experienced user.



NOTICE!

All figures in this chapter are showing examples using the management hyper terminal screen.

8.1 Host-side Parameters

Foreword: SCSI Channel, SCSI ID, and LUN

A SCSI channel (SCSI bus) can connect up to 15 devices (not including the controller itself) when the Wide function is enabled (16-bit SCSI). It can connect up to 7 devices (not including the controller itself) when the Wide function is disabled (8-bit SCSI). Each device has one unique SCSI ID. Two devices contending for the same SCSI ID are not allowed.




Fig. 8.1 SCSI ID/LUNs

Figure 8.1 shows this idea. If you file a document into a cabinet, you must put the document into one of the drawers. From a SCSI's point of view, a SCSI ID is like a cabinet, and the drawers are the LUNs (Logical units). Each SCSI ID enables up to 32 LUNs. Data can be stored into one of the LUNs of the SCSI ID. Most SCSI host adapters treat a LUN like another SCSI device.

8.1.1 Maximum Concurrent Host LUN Connection ("Nexus" in SCSI)

The configuration option adjusts the internal resources for use with a number of current host nexus. If there are four host computers (A, B, C, and D) accessing the array through four host IDs/LUNs (ID 0, 1, 2 and 3), host A through ID 0 (one nexus), host B through ID 1 (one nexus), host C through ID 2 (one nexus) and host D through ID 3 (one nexus) - all queued in the cache - that is called 4 nexus. If there are I/Os in the cache through four different nexus, and another host I/O comes down with a nexus different than the four in the cache (for example, host A access ID 3), the controller will return "busy." Note that it is "concurrent" nexus; if the cache is cleared up, it will accept four different nexus again. Many I/Os can be accessed via the same nexus.

8.1.3 Maximum Queued I/O Count

This function allows you to configure the maximum queued I/O count the controller can receive from the host computer.

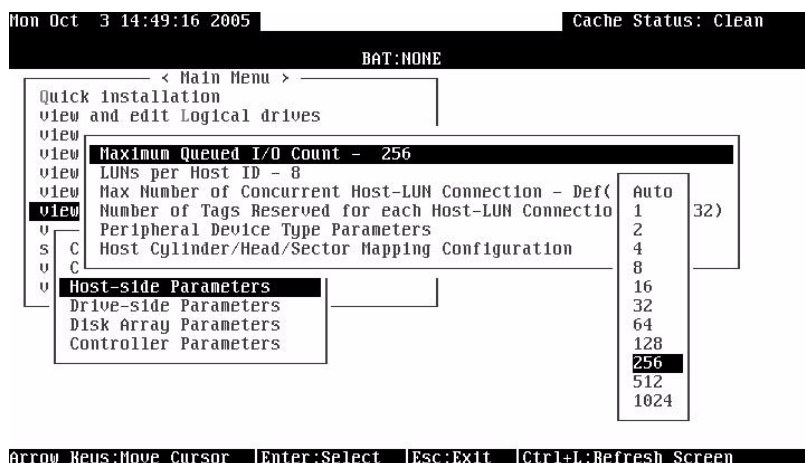


Fig. 8.4 Maximum Queued I/O Count

Choose “Host-side Parameters,” then press [ENTER]. Choose “Maximum Queued I/O Count,” then press [ENTER]. A list of available selections will appear. Move the cursor bar to an item, then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

The controller supports the following Host-side configurations:

“Maximum Queued I/O Count,” “LUNs per Host ID,” “Num of Host-LUN Connect,” “Tag per Host-LUN Connect,” “Peripheral Dev Type Parameters,” and “Cyl/Head/Sector Mapping Config.”

8.1.4 LUNs per Host ID

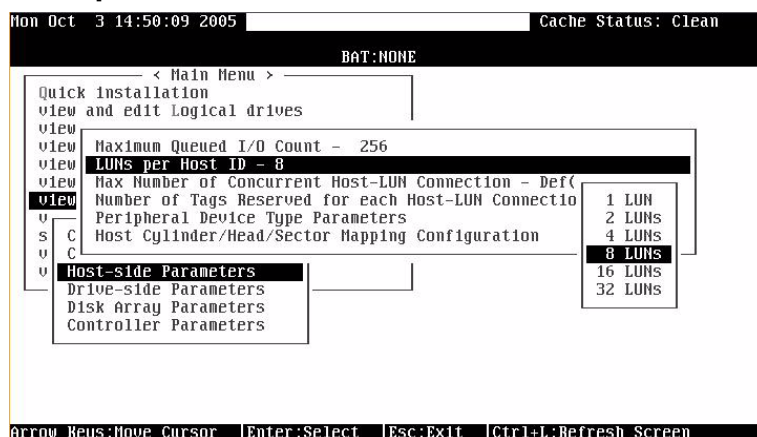


Fig. 8.5 LUNs per Host ID

Choose “LUNs per Host ID,” then press [ENTER]. A list of selections will appear. Move the cursor bar to an item, then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

8.1.5

LUN Applicability

If no logical drive has been created and mapped to a host LUN, and the RAID controller is the only device connected to the host SCSI card, usually the operating system will not load the driver for the host adapter. If the driver is not loaded, the host computer will not be able to use the in-band utility to communicate with the RAID controller. This is often the case when users want to start configuring a RAID using management software from the host. It will be necessary to configure the "Peripheral Device Type" setting for the host to communicate with the controller. If the "LUN-0's only" is selected, only LUN-0 of the host ID will appear as a device with the user-defined peripheral device type. If "all undefined LUNs" is selected, each LUN in that host ID will appear as a device with the user-defined peripheral device type. Different "LUN applicability" selections are available: "Device Type" selection, "Device Qualifier Support," "Support Removable media," "LUN-0's only," and "All undefined LUNs." Please refer to *Section 8.1.8 Peripheral Device Type Parameters for Various Operating Systems* for details concerning various operating systems.

8.1.6

Peripheral Device Type

For connection without a pre-configured logical unit and Ethernet link to a host, the in-band SCSI protocol can be used in order for the host to "see" the RAID subsystem. Please refer to the reference table below. You will need to make adjustments in the following submenus: Peripheral Device Type, Peripheral Device Qualifier, Device Support for Removable Media, and LUN Application.

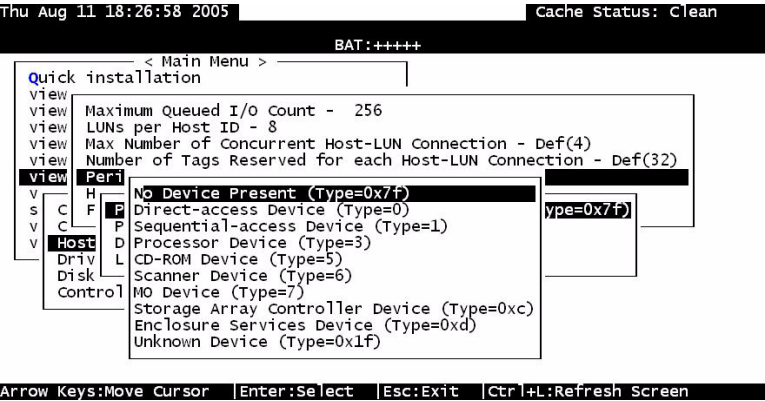


Fig. 8.6 Peripheral Device Type

8.1.7

In-band (SCSI channel)

External devices (including a RAID subsystem; from the view of operation on an application server or management PC) require communication links with a management computer for device monitoring and administration. In addition to the regular RS-232C or Ethernet connection, in-band SCSI can serve as an alternative means of management communications. In-band SCSI translates the original configuration commands into standard SCSI commands. These SCSI commands are then sent to and received by the controller over the existing host links, either SCSI or Fibre.

8.1.8 Peripheral Device Type Parameters for Various Operating Systems



NOTICE!

There is no need to configure the Peripheral Device setting if you are trying to manage a RAID subsystem from a RAIDWatch station through an Ethernet connection (to the Bosch subsystem's Ethernet port). An Ethernet connection to RAID uses TCP/IP as the communication protocol.

With an in-band connection, a host computer cannot “see” a RAID controller UNLESS the following have been configured:

1. A logical unit has been created.
or
At least one logical unit is mapped to a host ID or LUN via the RS-232/LCD keypad interface.
2. Host bus connection is established.
or
The RAID subsystem/controller is configured to appear as a peripheral device on the channel bus connected to a host computer.

With a brand new array, there is no association between disk drive configurations and the logical ID/LUN presentations on the host bus. If users want to start configuring a RAID system from an application server before any RAID configuration is made, the host will not be able to “see” the RAID subsystem. In order for a host to “see” the subsystem, it will be necessary to define the controller as a peripheral device.

Different host operating systems require different adjustments. See the tables below to find the proper settings for your host operating system. References to “Peripheral Device Qualifier” and “Device Support for Removable Media” are also included.

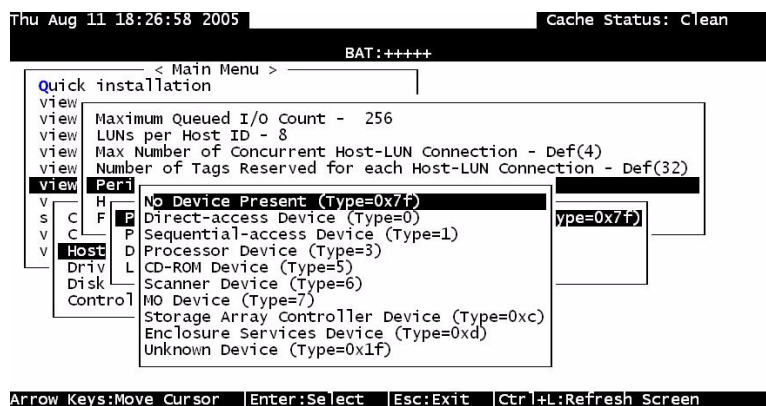


Fig. 8.7 Peripheral Device Type Parameters

Operating System	Peripheral Device Type	Peripheral Device Qualifier	Device Support for Removable Media	LUN Applicability
------------------	------------------------	-----------------------------	------------------------------------	-------------------

Table 8.1 Peripheral Device Type Parameters

Windows 2000/2003	0xd	Connected	Either is okay	LUN-0's
Solaris™ 8/9 (x86 and SPARC)	0xd	Connected	Either is okay	LUN-0's
Linux RedHat 8/9; SuSE 8/9	0xd	Connected	Either is okay	LUN-0's

Table 8.1 Peripheral Device Type Parameters

Device Type	Setting
Enclosure Service Device	0xd
No Device Present	0x7f
Direct-access Device	0
Sequential-access Device	1
Processor Type	3
CD-ROM Device	5
Scanner Device	6
MO Device	7
Storage Array Controller Device	0xC
Unknown Device	0x1f

Table 8.2 Peripheral Device Type Settings

8.1.9 Cylinder/Head/Sector Mapping

Drive capacity is decided by the number of blocks. For some operating systems (Sun Solaris, for example) the capacity of a drive is determined by the cylinder/head/sector count. For Sun Solaris, the cylinder cannot exceed 65535; choose "cylinder<65535," then the controller will automatically adjust the head/sector count for your OS to read the correct drive capacity. Please refer to the related documents provided with your operating system for more information.

Cylinder, Head, and Sector counts are selectable from the configuration menus shown below. To avoid any difficulties with a Sun Solaris configuration, the values listed below can be applied.

Capacity	Cylinder	Head	Sector
< 64 GB	variable	64	32
64 - 128 GB	variable	64	64
128 - 256 GB	variable	127	64
256 - 512 GB	variable	127	127
512 GB - 1 TB	variable	255	127

Table 8.3 Cylinder/Head/Sector Mapping under Sun Solaris

Older Solaris versions do not support drive capacities larger than 1 terabyte.

Solaris 10 now supports array capacity larger than 1TB. Set the values to the values listed in the table below:

Capacity	Cylinder	Head	Sector
>1TB	<65536	255	variable
		variable	255

Table 8.4 Cylinder/Head/Sector Mapping under Sun Solaris 10

Selecting Sector Ranges

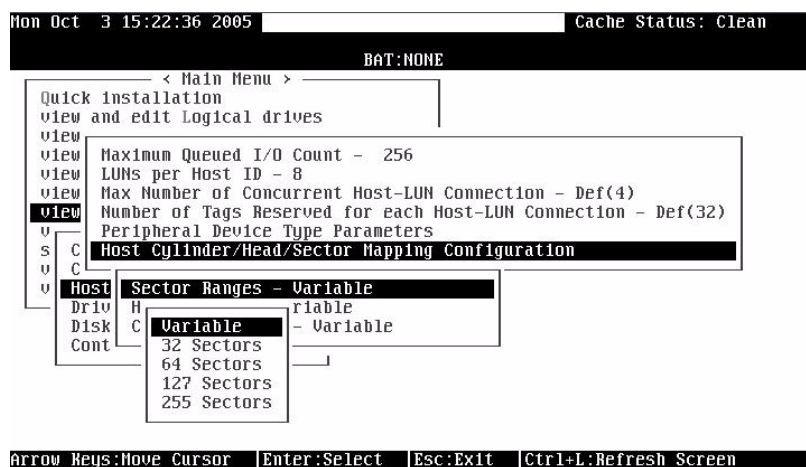


Fig. 8.8 Selecting Sector Ranges

Selecting Head Ranges

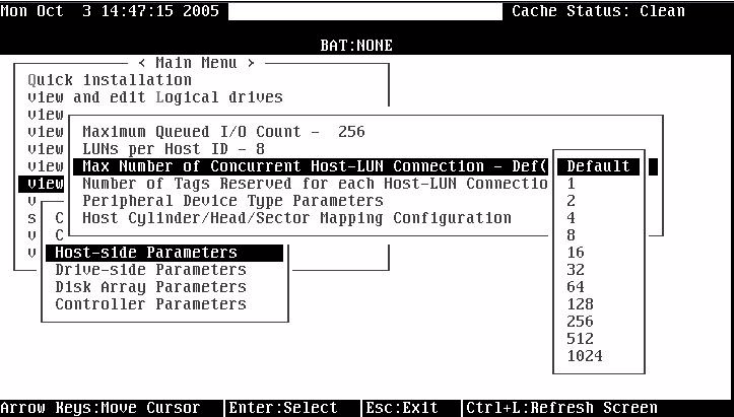


Fig. 8.9 Selecting Head Ranges

Selecting Cylinder Ranges

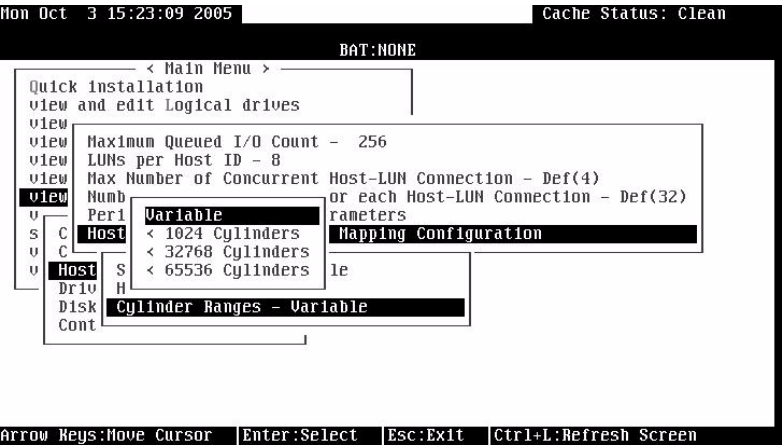
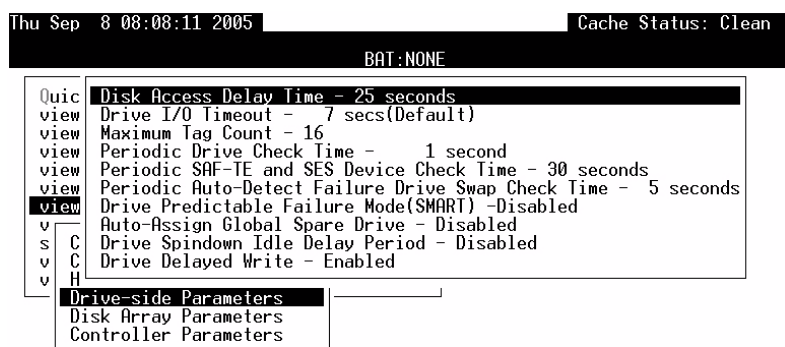


Fig. 8.10 Selecting Cylinder Ranges

8.2 Drive-side Parameters:



Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 8.11 Drive Side Parameters Screen

Choose “Drive-side Parameters,” then press [ENTER]. The Drive-side parameters menu displays.

8.2.1 Disk Access Delay Time

This feature sets the delay time before the subsystem tries to access the hard drives after power-on. Default may vary 15 seconds to 30 seconds, and from one model to another. This parameter can be adjusted to fit the spin-up speed of different models of disk drives installed in your subsystem.

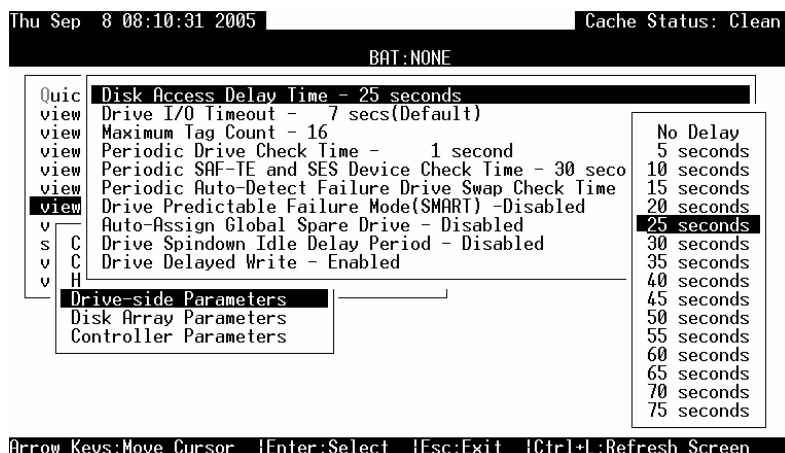
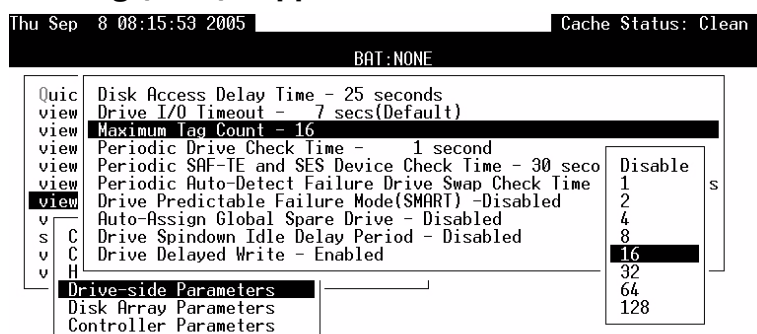


Fig. 8.12 Selecting Disk Access Delay Time

Choose “Disk Access Delay Time,” then press [ENTER]. A list of selections displays. Move the cursor bar to a selection, then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

8.2.3

Maximum Tag Count: Tag Command Queuing (TCQ) and Native Command Queuing (NCQ) Support

Arrow Keys: Move Cursor | Enter: Select | Esc: Exit | Ctrl+L: Refresh Screen

Fig. 8.14 Selecting Maximum Tag Count

This sub-menu facilitates the support for both Tagged Command Queuing (TCQ) and Native Command Queuing (NCQ). TCQ is a traditional feature on SCSI or Fibre disk drives, while NCQ is recently implemented with SATA disk drives. The queuing feature requires the support of both host adapters and hard disk drives. Command queuing can intelligently reorder host requests to streamline random accesses for IOPS/multi-user environments.

Bosch's subsystems support tag command queuing with an adjustable maximum tag count from 1 to 128. The default setting is "Enabled" with a maximum tag count of 32 (SCSI or Fibre drives) or 16 (default for SATA drives). Choose "Maximum Tag Count", then press [ENTER]. A list of available tag count numbers displays. Move the cursor bar to a number, then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

NOTICE!

- Every time you change this setting, you must reset the controller for the changes to take effect.
- Disabling Tag Command Queuing will disable the hard drives' built-in cache for Write-Back operation.
- The following options are categorized as related to array maintenance and data integrity:
 - Periodic Drive Check Time
 - Periodic Auto-Detect Failure Drive Swap Check Time
 - Auto-Assign Global Spare Drive
 Details of these options see *Section 10 Data Integrity*.
- Another option is associated with disk drive SMART support. Details see *Section 12 S.M.A.R.T. Configuration*.
 - Drive Predictable Failure Mode (SMART)



9 Enclosure Management

This chapter discusses the configuration options related to enclosure monitoring. Each function is given a brief explanation as well as a configuration sample. Terminal screens will be used in the configuration samples. Some of the operations require basic knowledge of RAID technology and are only recommended for an experienced user.



NOTICE!

All figures in this chapter are showing examples using the management hyper terminal screen.

9.1 Enclosure Monitoring

9.1.1 Enclosure Devices

Event Triggered Operations

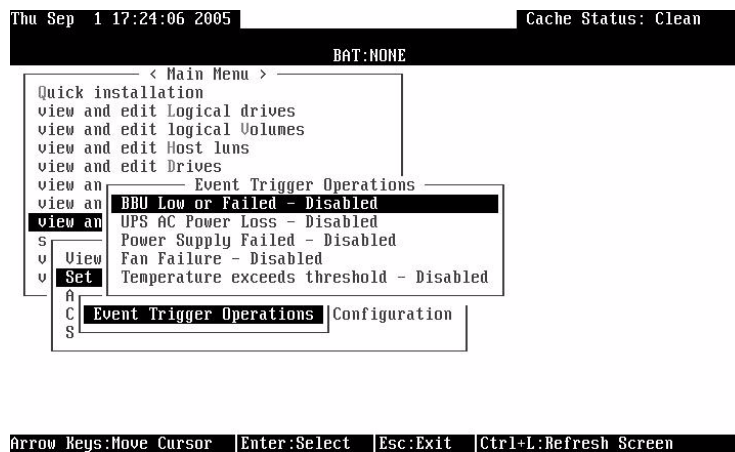


Fig. 9.1 Event Triggered Operations

Select “View and Edit Peripheral Devices” on the Main Menu and press [ENTER]. Choose “Set Peripheral Device Entry”, press [ENTER], then select “Event Trigger Operations” by pressing [ENTER]. The event trigger menu displays.

- The Operations

To reduce the chance of data loss, the controller/subsystem automatically commences the following actions to prevent loss of data:

 - a. Switches its caching mode from write-back to write-through
 - b. Flushes all cached data
- The Trigger

The mode switching and cache flush operations can be triggered by the detection of the following conditions:

 - a. BBU low or failed:

If battery fails or is under-charge, the unfinished writes cannot be supported if power outage occurs.
 - b. Power supply failure
 - c. Fan failure
 - d. Temperature exceeds threshold

If critical conditions occur, such as the enclosure component failure, chance of system downtime will increase and it is best to temporarily disable write-back caching.

When enabled, each of the above conditions forces the controller/subsystem to adopt the write-through caching mode. When the fault condition is corrected, the controller/subsystem automatically restores the previous caching mode.



Fig. 9.2 Controller Write-Through Alert

Note that the temperature thresholds refer to those set for both sensors on the RAID controller boards and those placed within the subsystem enclosure. In terms of the controller temperature, board 1 refers to the main circuit board and board 2 refers to the second-level I/O board or the daughter card. If any of the threshold values set for any sensor is exceeded, the trigger automatically applies.



NOTICE!

If a battery is not installed in your RAID subsystem, the “BBU Low or Failed “ option should be disabled.

9.1.2

Auto Shutdown: Elevated Temperature

System components can be damaged if operated under elevated temperature. You can configure the time periods between the detection of exceeded thresholds and the controller's commencing an automatic shutdown.

The shutdown does not electrically disconnect the subsystem. When shutdown is commenced, the subsystem stops responding to I/O requests and flushes all cached writes in its memory. During that time, system administrators should have been notified of the condition and have begun restoring proper cooling of the subsystem.

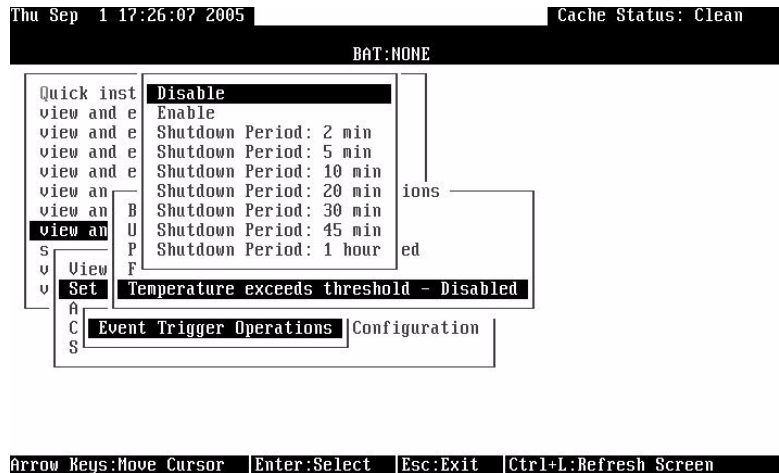


Fig. 9.3 Set Auto-Shutdown Period Time

Select “View and Edit Peripheral Devices” on the Main Menu and press [ENTER]. Choose “Set Peripheral Device Entry” and “Event Trigger Option” by pressing [ENTER]. The auto-shutdown options display.

Select a configurable time span between the detection of exceeded temperature and the controller's commencing an automatic shutdown. Extended operation under critical conditions like elevated temperature greatly reduces system efficiency and will eventually cause component failure.

Dynamic Caching Mode Switch

- Controller Auto-Shutdown - Event Triggered Operations

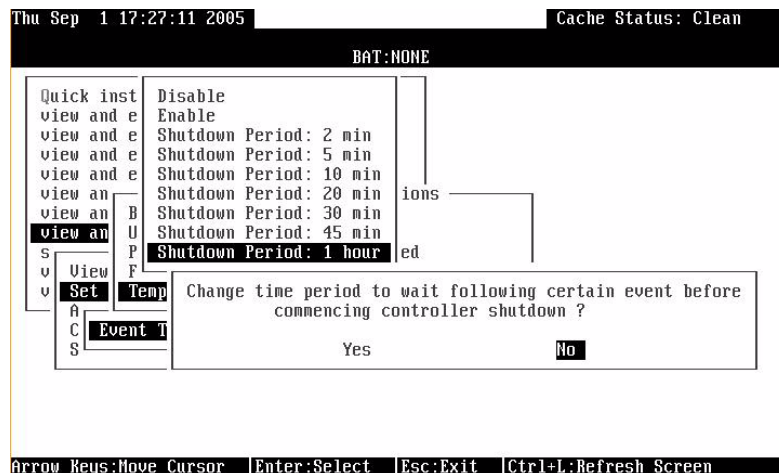


Fig. 9.4 Change Controller Shutdown Period

Select “View and Edit Peripheral Devices” on the Main Menu and press [ENTER]. Choose “Set Peripheral Device Entry” and “Event Trigger Option” by pressing [ENTER]. The auto-shutdown option displays.

Select a configurable time span between the detection of exceeded temperature and the controller's commencing an automatic shutdown. Extended operation under critical conditions like elevated temperature greatly reduces system efficiency and may cause component failure.

9.1.3

RAID Controller Self-monitoring

Open your PC Terminal Emulation screen. Enter the Main Menu and select “View and Edit Peripheral Devices.” Use the arrow keys to scroll down and select “Controller Peripheral Device Configuration,” “View Peripheral Device Status,” and then press [ENTER].

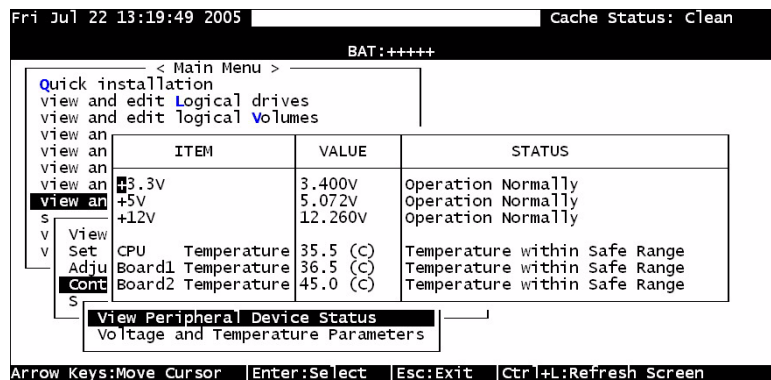


Fig. 9.5 View Peripheral Device Status

The controller operation status displays. The enclosure fan status is shown in the *Section 9.2.2 Enclosure Devices Status: View Peripheral Device Status..*

Changing Monitoring Thresholds

Open your PC Terminal Emulation utility. Enter the Main Menu and select “View and Edit Peripheral Devices.” Use the arrow keys to scroll down and select “Controller Peripheral Device Configuration,” “Voltage and Temperature Parameters,” and confirm by pressing [ENTER].

Note that it is not recommended to change the threshold values unless you need to coordinate the RAID controller's values with that of your RAID enclosure. If a value exceeding the safety range is entered, an error message will prompt and the new parameter will be ignored. For example, if the controller operates in a system enclosure where the upper limit on ambient temperature is relatively high or low, adjusting the default thresholds can coordinate the controller status monitoring with that of the system enclosure.

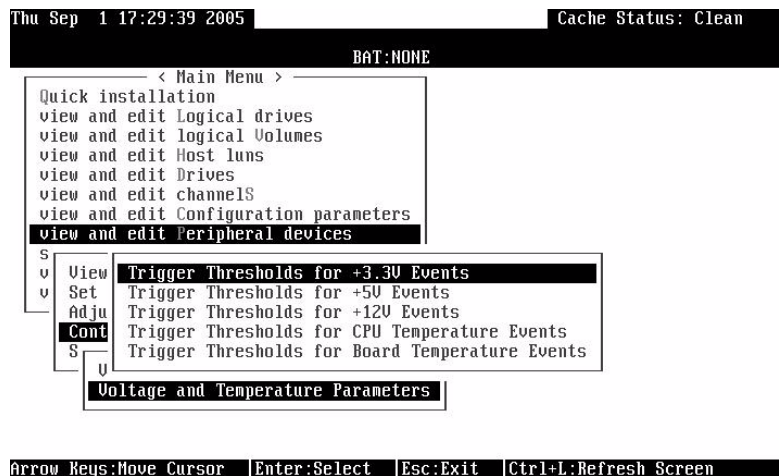


Fig. 9.6 Set Trigger Thresholds

Scroll down and select an item to configure.

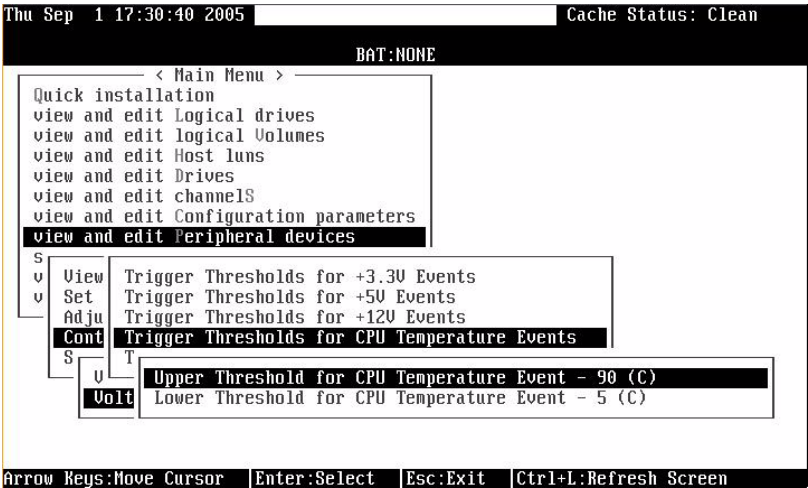


Fig. 9.7 Change Threshold for CPU Temperature

Select an item, such as “Trigger Thresholds for CPU Temperature Events.” Press [ENTER] and a list of selections will appear. You can change the upper or lower threshold values. Press [ENTER] to confirm.

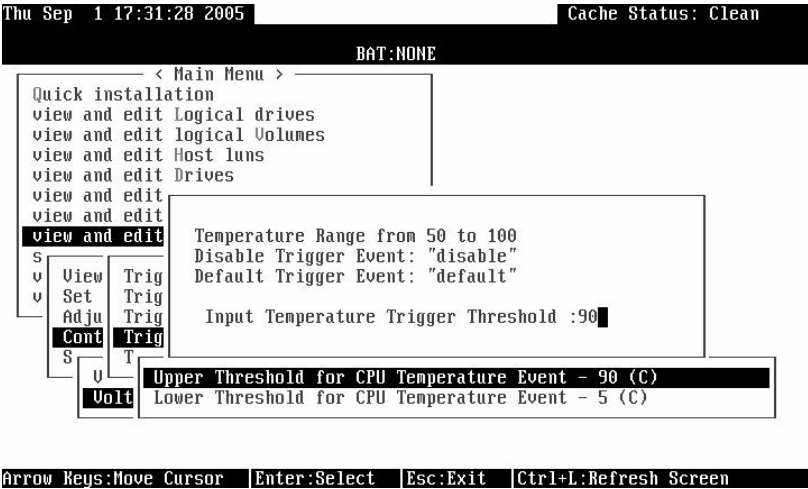


Fig. 9.8 Enter Trigger Threshold Value

A configuration window will prompt. Enter any value within the safety range. Values exceeding the safety range will be rejected by controller firmware.




Fig. 9.9 Warning Message for Out of Range

Follow the same method to modify other threshold parameters.

9.2 SAF-TE and SES Enclosure Monitoring

9.2.1 Periodic SAF-TE and SES Device Check Time

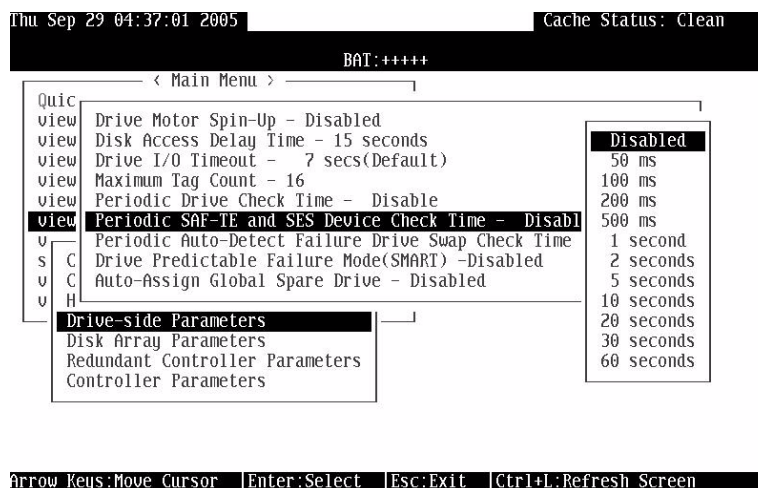


Fig. 9.10 Set SAF-TE and SES Check Time

If there are remote components in your RAID enclosure being controlled via SAF-TE/S.E.S. devices, use this function to select how often the controller will check the status of these devices. Choose “Periodic SAF-TE and SES Device Check Time”; then press [ENTER]. Move the cursor to the desired interval, then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

9.2.2 Enclosure Devices Status: View Peripheral Device Status

Select “View and edit Peripheral Devices” on the Main Menu and press [ENTER]. Choose “View Peripheral Device Status,” then press [ENTER] again. The device list displays. Below is a list of peripheral devices (enclosure modules) supported by enclosure monitoring. Monitoring of device status depends on enclosure implementation and is accessed through different interfaces, e.g., SAF-TE, S.E.S., or I2C serial bus.

1. Device type
2. Enclosure descriptor
3. Power supply
4. Cooling element
5. Temperature Sensors
6. Audible alarm
7. Enclosure services controller electronics

Below is a screen display showing enclosure devices through an I2C serial bus:

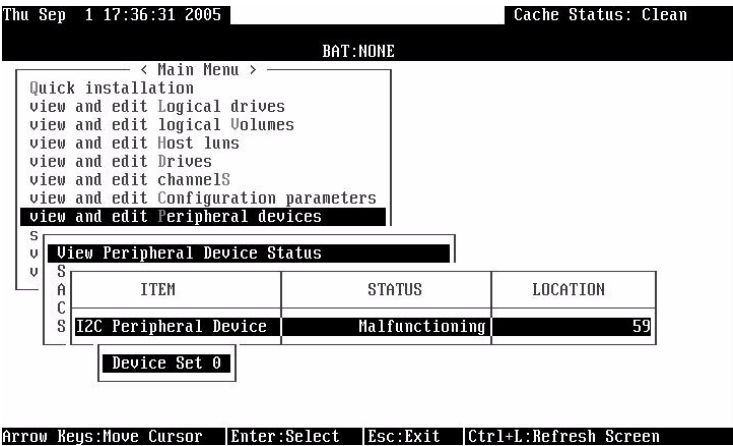


Fig. 9.11 I2C Device Information

Select the device interface and then press [ENTER] on a connected module to examine its operating status as shown below.

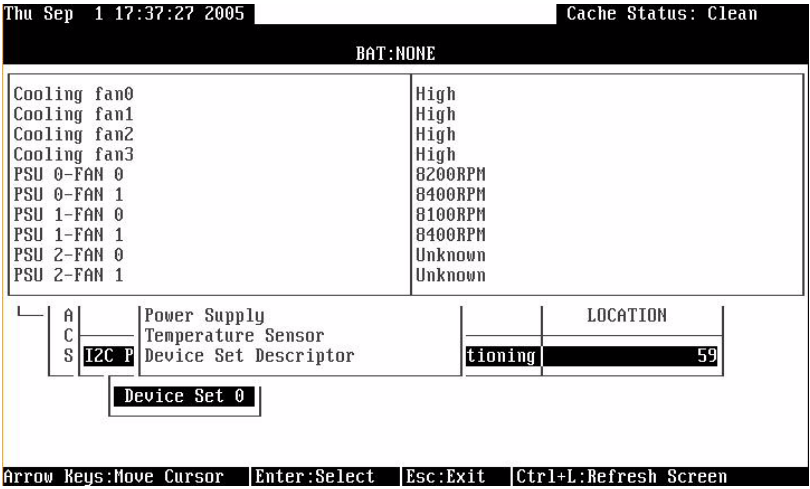


Fig. 9.12 Display I2C Device Information

10 Data Integrity

This chapter discusses various firmware mechanisms that help to secure the integrity of array data.

No system is completely safe from the threat of hardware faults. For example, although the chance of occurrence is considerably low, the occurrences of bad blocks on two (RAID 5) hard drives can fail a whole data set. When properly configured, the functions below help to minimize the chance of data loss:

1. Event Triggered Operations (refer to *Section 9 Enclosure Management*)
2. <XREF>Failed Drive Detection
3. <XREF>Scheduled Maintenance
4. <XREF>Regenerate Logical Drive Parity
5. <XREF>Rebuild Priority
6. <XREF>Verification on Writes



NOTICE!

- Some of the configuration options may not be available to all sub-revisions of firmware.
- All figures in this chapter are showing examples using the management hyper terminal screen.

10.1 Failed Drive Detection

10.1.1.1 Periodic Drive Check Time

The “Periodic Drive Check Time” is the time interval for the controller to check all disk drives that were on the drive bus at controller startup (a list of all the drives that were detected can be seen under “View and Edit Drives”).

The default value is “1 second.” If set to “Disabled,” it means that if a drive is removed from the bus, the subsystem will not be able to know – so long as no host accesses that drive.

Changing the check time to any other value allows the controller to check – at the selected time interval – all of the drives that are listed under “View and Edit Drives.” If any drive is then removed, the controller will be able to know – even if no host accesses that drive.

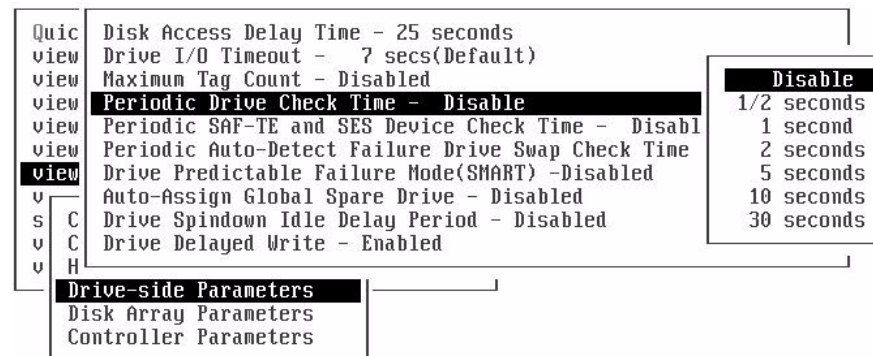


Fig. 10.1 Set for Periodic Drive Check Time

Choose “Periodic Drive Check Time,” then press [ENTER]. Move the cursor to the desired interval, then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

10.1.2

Periodic Auto-Detect Failure Drive Swap Check Time

The “Drive-Swap Check Time” is the interval at which the controller checks to see if a failed drive has been swapped. When a logical drive’s member drive fails, the controller will detect the failed drive (at the selected time interval). Once the failed drive has been swapped with a drive that has the adequate capacity to rebuild the logical drive, the rebuild will begin automatically

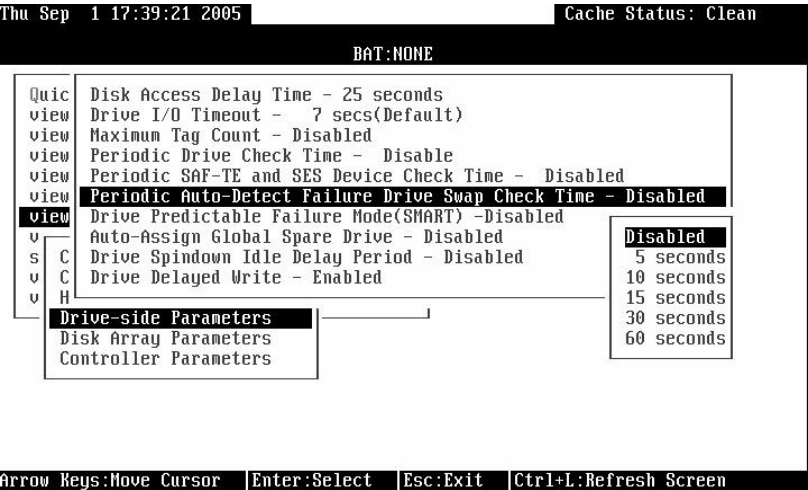


Fig. 10.2 Set for Auto-Detect Failure Drive Swap Check Time

Choose “Periodic Auto-Detect Failure Drive Swap Check Time”; then press [ENTER]. Move the cursor to the desired interval; then press [ENTER]. Choose Yes in the dialog box that follows to confirm the setting.

NOTICE!

The "Periodic Drive Check Time" is enabled by choosing a time value. The RAID controller will poll all connected drives through the controller’s drive channels at the assigned interval. Drive removal will be detected even if a host does not attempt to access data on that specific drive. If the "Periodic Drive Check Time" is set to "Disabled" (the default setting is "Disabled"), the controller will not be able to detect any drive removal that occurs after the controller initialization process. The controller will only be able to detect drive removal when host access is directed to the drive side.

10.1.3

Auto-Assign Global Spare Drive

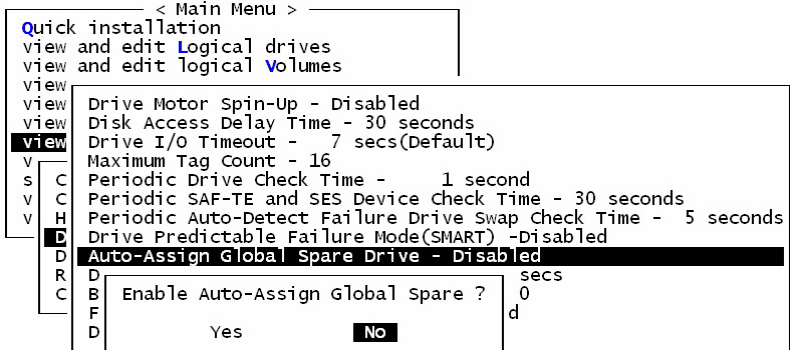


Fig. 10.3 Enable Auto-Assign Global Spare Drive

The “Auto-Assign” function automatically assigns any “new” drives that are not included in logical configurations as Global Spares.

The Fault Scenario:

Spare drives accelerate rebuild of a logical drive. In the example as described below, multiple faults can occur at the same time making the array exposed to the risk of data loss:

- There is only one Global Spare in a RAID subsystem.
- That Global Spare has been used to rebuild a logical drive.
- The failed drive is swapped out and replaced by a new one.
- Chances are system administrators forgot to configure the replacement drive as a spare.
- A member of another logical drive fails.
- The subsystem has no spare left. Performance decreases for a considerable portion of system resources has to be conducted to generate data from the remaining members of the logical drive. If yet another member fails in the logical drive, data is lost.

The chance of failing two drives increases when a failed drive in the array cannot be replaced immediately for the lack of spare drives.

The Function

If a drive has a capacity smaller or apparently larger than the members of configured arrays, the controller may avoid using it as a global spare.

Enable the function and reset the controller for the configuration to take effect.

10.2 Scheduled Maintenance

10.2.1 Task Scheduler

The Task Scheduler functionality allows Media Scans to be scheduled beginning at a specified start time and repeating at regular intervals defined by a configurable interval period. Each such schedule can be defined to operate on individual drives, all drives of a certain class, all member drives of a specified logical drive, or all member drives of all logical drives. UIs supported are RS232C terminal menus and RAIDWatch GUI manager.

Fri May 7 13:44:51 2004

Cache Status: Clean

BAT:NONE

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0	ECC3E8E	NA	RAID5	1673784	SHUTDOWN					7	B	8	1	0
<div>View drives Delete logical drive Partition logical drive logical drive Name Expand logical drive add drives reGenerate parity copy and replace drive Media scan Write policy</div>				112	GOOD					7	B	2	1	0
				000	GOOD					7	B	2	1	0
				000	GOOD					7	B	3	1	0
7			NONE											

Arrow Keys:Move Cursor

Enter:Select

Esc:Exit

Ctrl+L:Refresh Screen

Fig. 10.4 Media Scan Function

The scheduler is accessed through the “View and Edit Logical Drives” menu by selecting a configured array to display the list of array-related functions. Select “Media Scan” by pressing [ENTER].

Media Scan Priority - Normal
Iteration Count - Single Time
Task Scheduler

Fig. 10.5 Task Scheduler Function

Select “Task Scheduler” by pressing [ENTER].

Creating a New Schedule

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL
0	ECC3E8E	NA	RAID5	1673784					7	B	8	1	0
	Media Scan Priority - Normal								7	B	2	1	0
	Iteration Count - Single Time												
	Task Scheduler								7	B	2	1	0
3	No Media Scan Task Schedule - Add a New Task Schedule ?												
4	<div>Yes</div> <div>No</div>												

Fig. 10.6 Enable Task Scheduler

If there is no preset schedule, a confirm box as shown above will prompt.

Adding or Deleting a Schedule

Media Scan Priority - Normal				7	B	2	1	0	
Iteration Count - Single Time				7	B	2	1	0	
Task Scheduler									
3	Idx	Start Time and Date		Period			Exec on Init		
4	0	Fri May 7 14:09:00 2004		10 hours			Priority		
5	View Schedule Information								
6	Add New Schedule								
	Delete Schedule								

Fig. 10.7 Adding or Deleting Task Scheduler

Press [ENTER] on an existing schedule to display the configuration options. You may choose to check information of a task schedule, to create a new schedule, or to remove a configured schedule.

10.2.2

Task Scheduler Options

To configure a task schedule, you have the option to do any of the following:

Set the following values/intervals:

Execute on Controller Initialization

Media Scan Priority - Normal
Iteration Count - Single Time
Task Scheduler
Execute on Controller Initialization - YES
Start time and date - Not Set
Execution Period - Not Set
Media Scan Mode - Concurrence
Media Scan Priority - Normal
Select logical drives

Fig. 10.8 Set for Task Scheduler Function

This option determines whether Media Scan is automatically conducted whenever the RAID system is reset or powered on.

Start Time and Date

Execute on Controller Initialization - YES	1	0
Start time and date - Fri May 7 14:09:00 2004		
E M M S	Set Start Time and Date [MMDDhhmmYYYY] : ■	

Fig. 10.9 Set for Date on Task Scheduler Function

Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

Execution Period

Execute on Controller Initialization - YES	Schedule Period - YES
Start time and date - Fri May 7 14:09:00 2004	
Execution Period -	Execution Once
Media Scan Mode - Concurrence	1 secs
Media Scan Priority - Normal	2 secs
Select logical drives	3 secs
	⋮
	10 days
	11 days
	12 days
	13 days
	2 weeks
	3 weeks

Fig. 10.10 Set for Time on Task Scheduler Function

The scheduler memorizes the date and the time the actions are to be executed. Select one of the following:

- If the action is intended to be executed for one time only, select “Execution Once.”

- In the case of a periodic action, the action is executed at the specified “start time,” and then re-enacted at the time interval indicated in the execution period so as to be executed again later. The selectable interval ranges from one second to several weeks.

Media Scan Mode




Fig. 10.11 Set for Media Scan Mode

If the maintenance schedule includes more than one logical drive, the scan can be performed simultaneously or separately on one disk at a time in a sequential order.

Media Scan Priority




Fig. 10.12 Set for Media Scan Priority

The scan priority determines how much of the system’s resources will be consumed to perform the scheduled task. Select “Low” for better array performance and longer time to complete the media scan. Higher priority allows higher scan performance at the cost of reduced array performance.

Select Logical Drives

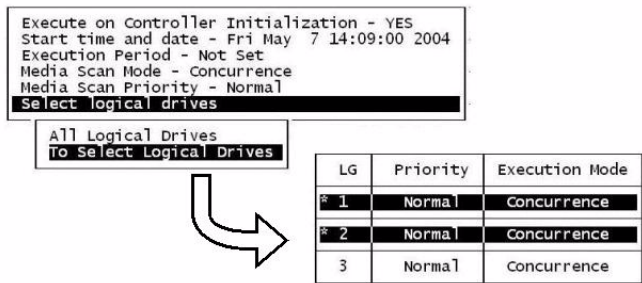


Fig. 10.13 Set for Logical Drive

Press [ENTER] on “Select Logical Drives” to bring out a sub-menu. From there you may include all configured arrays or press [ENTER] on “To Select Logical Drives” to select one or more specific logical drive(s). Logical drives can be tagged for inclusion by positioning the cursor bar on the logical drive and then pressing [ENTER]. An asterisk (*) mark will appear on the selected physical drive(s). To deselect the drive, press [ENTER] again on the selected drive. The “*” mark will disappear. Use the same method to select more logical drives if required. When selection is done, press [ESC] to continue.

10.2.3

Confirming the Creation of a Task Schedule

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#
0	ECC3E8E	NA	RAID5	1673784	SHUTDOWN				7	B	8	1	
	Media Scan Priority - Normal									7	B	2	1
	Iteration Count - Single Time									7	B	2	1
	Task Scheduler									7	B	2	1
3	Start Time and Date : Fri May 7 14:09:00 2004												1
4	Schedule Period : 10 hours												
5	Execution Mode : Sequence												
6	Priority : Low												
7	Logical Drive Count : 1												
6	Create Task Schedule ?												
7	Yes No												

Fig. 10.14 Enable Task Scheduler Function

When finished with setting the scheduler options, press [ESC] to display a confirm box. Verify all information in the box before choosing “Yes” to confirm and to complete the configuration process.

10.3 Regenerating Logical Drive Parity

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME	
P0	4295529B	NA	RAID5	476000	GOOD						4	B	3	0	0
	Execute Regenerate Logical Drive Parity														
	Overwrite Inconsistent Parity - Enabled														
	Generate Check Parity Error Event - Enabled														
3			NONE												
4			NONE												
5			NONE												
6			NONE												
7			NONE												

Fig. 10.15 Regenerate Logical Drive Parity

Parity regeneration is a function manually performed onto RAID-1/3/5 arrays to determine whether inconsistency has occurred with data parity.

You may perform the parity check directly without changing the two options below, or set preferred options and then press [ENTER] on “Execute Regenerate Logical Drive Parity” to begin the operation.

10.3.1 Overwrite Inconsistent Parity


Default is “enabled.”

If an array’s data parity is seriously damaged, restoring parity data by regenerating and overwriting the original data may cause data loss. Disable this option if you suspect parity data has been seriously corrupted.

10.3.2 Generate Check Parity Error Event

Default is “enabled.”

When enabled, parity inconsistency will be reported as system events.



NOTICE!

If a regenerating process is stopped by a drive failure, the process cannot be restarted until the logical drive is successfully rebuilt by having its failed member replaced.

10.4 Disk Array Parameters

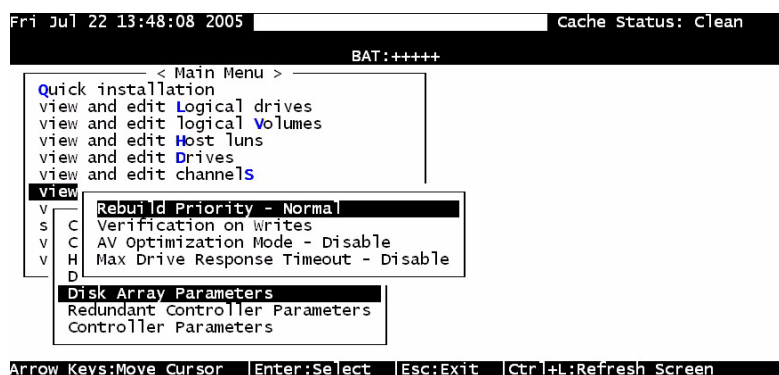


Fig. 10.16 Set for Disk Array Parameters

Select “View and Edit Configuration Parameters” on the Main Menu and press [ENTER]. Choose “Disk Array Parameters,” then press [ENTER] again. The Disk Array Parameters menu will appear.

10.4.1 Rebuild Priority

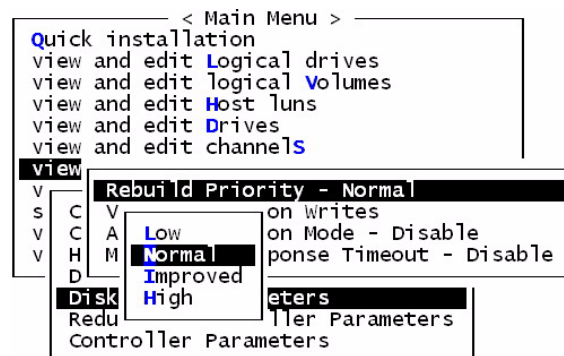


Fig. 10.17 Change Rebuild Priority

Choose “Rebuild Priority,” then press [ENTER]. A list of the priority selections (Low, Normal, Improved, or High) displays. Move the cursor bar to a selection, then press [ENTER].

10.4.2 Verification on Writes

Errors may occur when a hard drive writes data. To avoid the write error, the controller can force hard drives to verify written data.

There are three selectable methods:

- Verification on LD Initialization Writes
Performs Verify-after-Write when initializing a logical drive
- Verification on LD Rebuild Writes
Performs Verify-after-Write during the rebuild process
- Verification on LD Normal Drive Writes
Performs Verify-after-Write during normal I/Os

Each method can be enabled or disabled individually. Hard drives will perform Verify-after-Write according to the selected method.

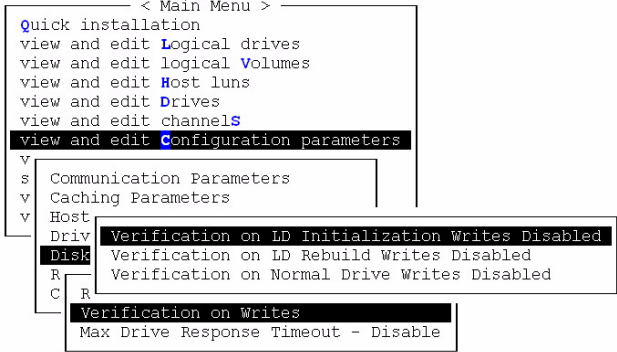


Fig. 10.18 Configure Verification on Write

Move the cursor bar to the desired item, then press [ENTER].

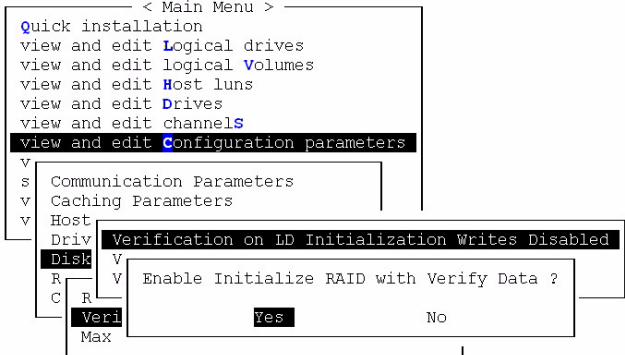


Fig. 10.19 Enable Verification on LD Initialization Writes

Choose Yes in the confirm box to enable or disable the function. Follow the same procedure to enable or disable each method.



NOTICE!
The “verification on Normal Drive Writes” method will affect the “write” performance of your RAID system.

11 Array Expansion

The array expansion functions allow you to expand storage capacity without the cost of buying new equipment. Expansion can be completed on-line while the system is serving host I/Os.

11.1 Overview

What is RAID Expansion and how does it work?

Before the invention of RAID Expansion, increasing the capacity of a RAID system meant backing up all data in the disk array, re-creating the disk array configuration with new drives, and then restoring data back into system.

RAID Expansion technology allows users to expand a logical drive by adding new drives, or replacing drive members with drives of larger capacity. Replacing is done by copying data from the original members onto larger drives; the smaller drives can then be replaced without powering down the system.

11.2 Notes on Expansion

11.2.1 Added Capacity

When a new drive is added to an existing logical drive, the capacity brought by the new drive appears as a new partition. For example, if you have 4 physical drives (36GB each) in a logical drive, and each drive's maximum capacity is used, the capacity of the logical drive will be 108GB. (One drive's capacity is used for parity, e.g., RAID 3). When a new 36GB drive is added, the capacity will be increased to 144GB in two separate partitions (one is 108GB and the other 36GB).

11.2.2 Size of the New Drive

A drive used for adding capacity should have the same or more capacity as other drives in the array.

11.2.3 Applicable Arrays

Expansion can only be performed on RAID 0, 1, 3 and 5 logical drives. Expansion cannot be performed on logical configurations that do not have parity, e.g., RAID or RAID 1.



NOTICE!

Expansion on RAID0 is not recommended, because the RAID0 array has no redundancy. Interruptions during the expansion process may cause unrecoverable data loss.

11.2.4

Interruption to the Process

Expansion should not be canceled or interrupted once begun. A manual restart should be conducted after the occurrence of a power failure or interruption of any kind.

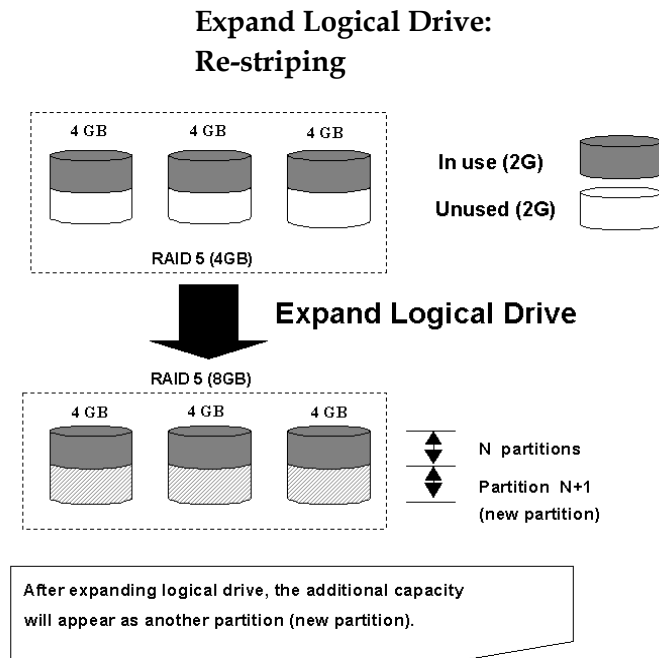


Fig. 11.1 Logical Drive Expansion

RAID levels supported: RAID 0, 1, 3 and 5

Expansion can be performed on logical drives or logical volumes under the following conditions:

1. There is unused capacity in a logical unit
2. Capacity is increased by using member drives of larger capacity

Data is recalculated and distributed to drive members or members of a logical volume. Upon the completion of the process, the added or the previously unused capacity will become a new partition. The new partition must be made available through host LUN mapping in order for a host adapter to recognize its presence.

11.3

Mode 1 Expansion:

Adding Drives to a Logical Drive

Use drives with the same capacity as the original drive members. Once completed, the added capacity will appear as another partition (new partition). Data is automatically re-stripped across the new and old members during the add-drive process. See the diagram below to get a clear idea:

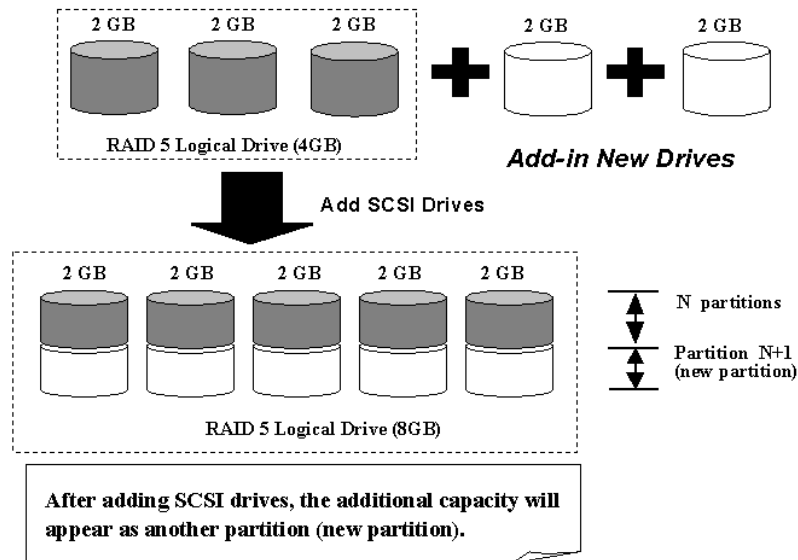


Fig. 11.2 Expansion by Adding Drive

RAID levels supported: RAID 0, 1, 3 and 5.

The new partition must be made available through a host ID/LUN.

Add Drive Procedure

First select from the Main Menu, “View and Edit Logical Drive,” and select a logical drive to add a new drive to. The drive selected for adding should have a capacity no less than the original member drives. If possible, use drives of the same capacity because all drives in the array are treated as though they have the capacity of the smallest member in the array.

NOTICE!

The following figures in this chapter are showing examples using the management hyper terminal screen.

Cache Status: Clean											
Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
0	P0	7559F503	NA	RAID5	1279	GOOD	R	5	0	0	
1				NONE							
2				NONE							
3				NONE							
4				NONE							
5				NONE							
6				NONE							
7				NONE							

Arrow Keys: Move Cursor | Enter: Select | Esc: Exit | Ctrl+L: Refresh Screen

Fig. 11.3 Add Drive on Logical Drive

Press [ENTER] to select a logical drive and choose “Add Drives” from the submenu. Proceed with confirming the selection.

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0	5DC80200	0	RAID6	4000	GOOD				7	8	4	0	0	
<div><div>view drives</div><div>Delete logical drive</div><div>Partition logical drive</div><div>logical drive Name</div><div>Expand logical drive</div><div>Migrate logical drive</div><div>add drives</div><div>reGenerate parity</div><div>copy and replace drive</div><div>Media scan</div><div>write policy</div></div>														
7			NONE											

Fig. 11.4 Choose Add Drives Function

Available drives will be listed. Select one or more drive(s) to add to the target logical drive by pressing [ENTER]. The selected drive will be indicated by an asterisk “*” mark.

LG		Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
0	2A	8	238469	150MB	NONE	FRMT DRV	
1	12	9	238469	150MB	NONE	FRMT DRV	
	view	10	238469	150MB	NONE	FRMT DRV	
	Dele	11	238469	150MB	NONE	FRMT DRV	
	Part						

Fig. 11.5 Select Drives

Press [ESC] to proceed and the notification will prompt.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
<	0	5DC80200	0	RAID6	9999	GOOD	R	3	0	0	
<	Adding Notification										
<	LG:0 Logical Drive NOTICE: Starting Add Drive Operation										
<	4			NONE							
<	5			NONE							
<	6			NONE							
<	7			NONE							

Fig. 11.6 Notice Starting Add Drive Operation

Press [ESC] again to cancel the notification prompt; a status bar will indicate the percentage of progress.

Q	LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
<	0	2E5B167A	0	RAID6	9999	GOOD	R	3	0	0	
<	Adding										
<	1										
<	2										
<	3										
<	4			NONE							
<	5			NONE							
<	6			NONE							
<	7			NONE							

Fig. 11.7 View the Percentage of Progress

Upon completion, there will appear a confirming notification. The capacity of the added drive will appear as an unused partition.

Adding Notification										
Add Drive to Logical Drive 0 Completed										

11.4 Mode 2 Expansion

Copy and Replace Drives with Drives of Larger Capacity

You may also expand your logical drives by copying and replacing all member drives with drives of higher capacity. Please refer to the diagram below for a better understanding. The existing data in the array is copied onto the new drives, and then the original members can be removed.

When all the member drives have been replaced, execute the “Expand Logical Drives” function to make use of the added capacity.

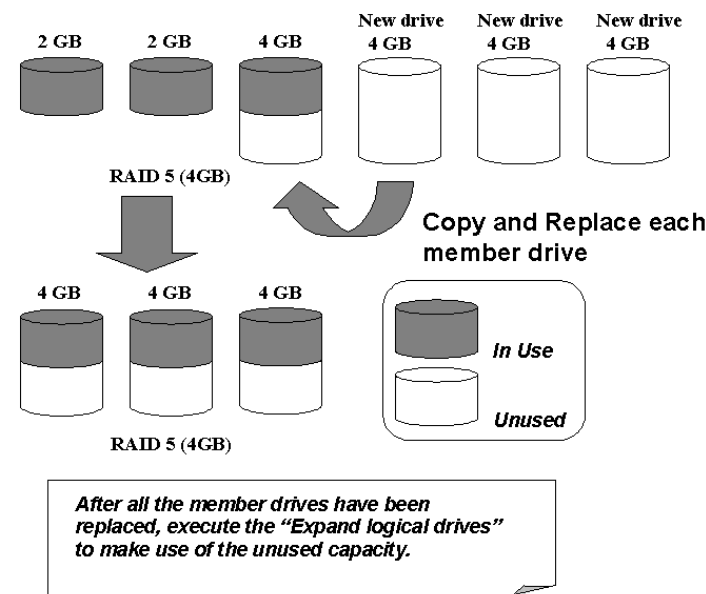


Fig. 11.9 Expansion by Copy & Replace

RAID levels supported: RAID 0, 1, 3 and 5

Copy and Replace Procedure

Select from Main Menu “View and Edit Logical Drives.” Select a target array, press [ENTER] and scroll down to choose “Copy and Replace Drive.” Press [ENTER] to proceed.

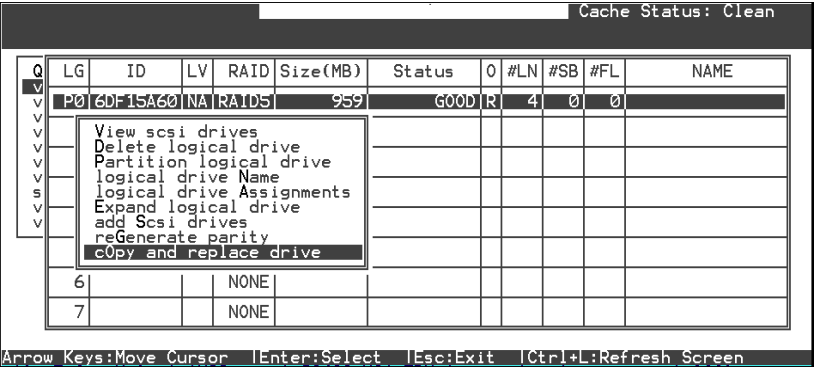


Fig. 11.10 Choose Copy & Replace

The array members will be listed. Select the member drive (the source drive) you want to replace with a larger one.

11.5 Making Use of the Added Capacity: Expand Logical Drive

In the following example, the logical drive is originally composed of three member drives and each member drive has the capacity of 1GB. “Copy and Replace” has been performed on the logical drive and each of its member drives has been replaced by a new drive with the capacity of 2GB. The next step is to perform “Expand Logical Drive” to utilize the additional capacity brought by the new drives.

- 1. Select “View and Edit Logical Drives” from the Main Menu and select the logical drive with its members copied and replaced.
- 2. Select “Expand Logical Drive” in the sub-menu and press [ENTER] to proceed. A confirming box will appear.
- 3. Proceed by pressing [ENTER] or entering any value no larger than the "maximum drive expand capacity" and press [ENTER].

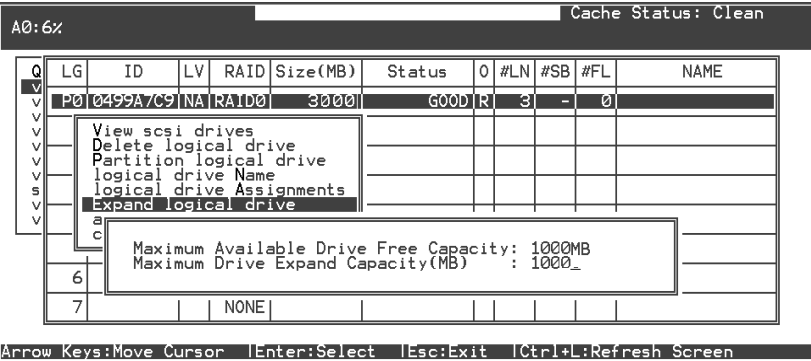


Fig. 11.16 Expand Logical Drive

Choose Yes to confirm and proceed.




Fig. 11.17 Choose Expand Logical Drive

Upon completion, you will be prompted by the notification message.




Fig. 11.18 Completed Expand Logical Drive

Press [ESC] to return to the previous menu screen.
The total capacity of logical drive has been expanded to 6GB.

Q V V V V V V S V V	LG	ID	LV	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
	P0	0499A/C9	NA	RAID0	6000	Good	R	3	-	0	
	1			NONE							
	2			NONE							
	3			NONE							
	4			NONE							
	5			NONE							
	6			NONE							
	7			NONE							

Fig. 11.19 View Logical Drive Information

11.6 Expand Logical Volume

To expand a logical volume, expand its logical drive member(s) and then perform “expand logical volume.”

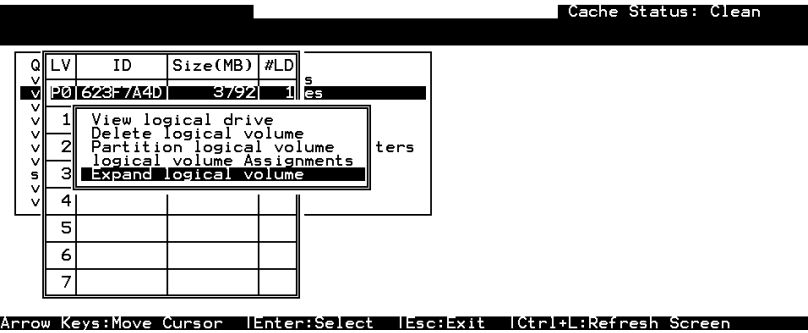


Fig. 11.20 Expand Logical Volume

When prompted by "Expand Logical Volume?", choose Yes to confirm and the process will be completed immediately.

11.7 Configuration Example: Volume Extension in Windows 2000®

Limitations When Using Windows 2000

1. This limitations apply only to the Windows NT Server or Windows 2000 Server Disk Management which includes the Extend Volume Set function; Windows NT Workstation does not support this feature. The Volume Set Expansion formats the new area without affecting existing files on the original volume.
2. The system drive (boot drive) of a Windows NT/2000 system cannot be expanded.
3. The drive to be expanded should be using the NTFS file system.

Example

The following example demonstrates the expansion of a 16988MB RAID 5 logical drive. The HyperTerminal emulation software that comes with Windows Server is used to connect to the RAID controller via RS-232C.

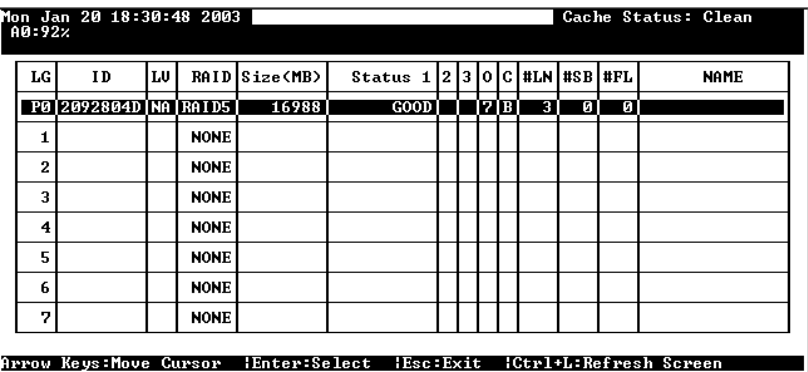


Fig. 11.21 View Logical Drive

You can view information about this drive in the Windows 2000 Server's Computer Management -> Storage -> Disk Management.

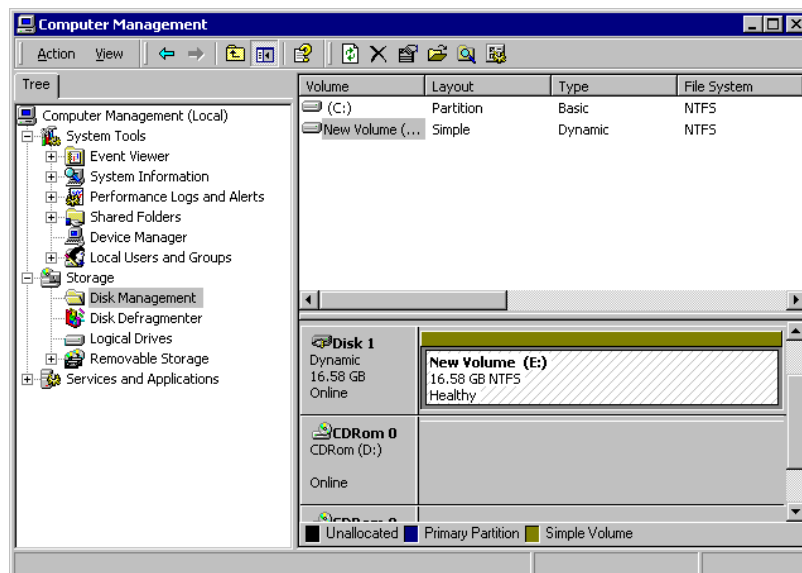


Fig. 11.22 View Disk Management on Windows

Place the cursor on Disk 1. Then right-click your mouse and select “Properties”. You will see that the total capacity for the Drive E: is about 16.5GB.

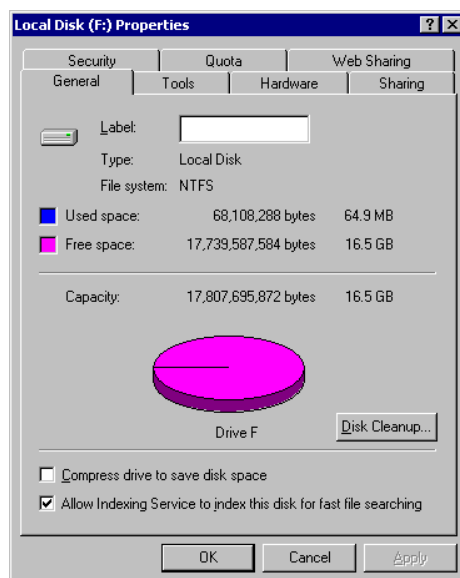


Fig. 11.23 View Disk Information

Follow the steps described in the previous section to "add" or "copy & replace" disk drives and perform “Logical Drive Expansion.”

Mon Jan 20 18:30:48 2003 Cache Status: Clean
A0:92%

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	2092804D	NA	RAID5	16988	GOOD						7	8	3	0
1					Adding									
2														
3														
4			NONE											
5			NONE											
6			NONE											
7			NONE											

92% Completed_

Arrow Keys:Move Cursor !Enter:Select !Esc:Exit !Ctrl+L:Refresh Screen

Fig. 11.24 View Logical Drive Expansion of progress

The 16.5GB logical drive has become a 25GB logical drive. Place the cursor on that logical drive, and then press [ENTER].

Mon Jan 20 18:32:34 2003 Cache Status: Clean

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	2092804D	NA	RAID5	25482	GOOD						7	8	4	0
1			NONE											
2			NONE											
3			NONE											
4			NONE											
5			NONE											
6			NONE											
7			NONE											

Arrow Keys:Move Cursor !Enter:Select !Esc:Exit !Ctrl+L:Refresh Screen

Fig. 11.25 Choose Logical Drive

From the menu, select "Partition Logical Drive." You will see that the 25GB logical drive is composed of a 17GB partition and an 8.4GB partition.

Mon Jan 20 18:33:39 2003 Cache Status: Clean

LG	ID	LU	RAID	Size(MB)	Partition	Offset(MB)	Size(MB)	NAME
P0	2092804D	NA	RAID5	25482	0	0	16988	
1			NONE		1	16988	8494	
2			NONE		2			
3			NONE		3			
4			NONE		4			
5			NONE		5			
6			NONE		6			
7			NONE		7			

Arrow Keys:Move Cursor !Enter:Select !Esc:Exit !Ctrl+L:Refresh Screen

Fig. 11.26 View Partition of Logical Drive

Follow the directions in *Section 5 LCD Keypad Operation* and *Section 7 Terminal Operation* to map the new partition to a host LUN. The new partition must be "mapped" to a host LUN in order for the HBA (host-bus adapter) to see it. Once you have mapped the partition, reboot your Windows server. The HBA should be able to detect an additional disk during the initialization process.

Return to Windows 2000 Server's Disk Management. There now exists a Disk 2 with 8.3GB of free space. You may use the "rescan disks" command to bring up the new drive.

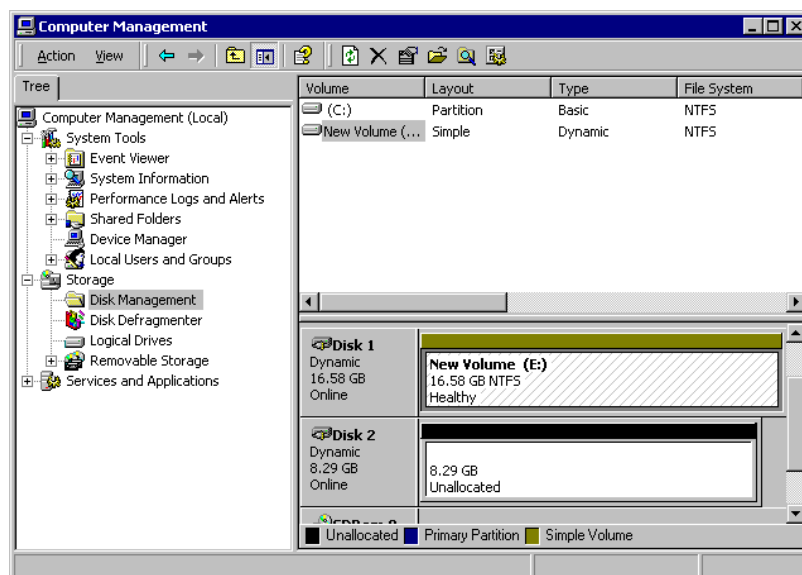


Fig. 11.27 View Disk information on Windows

Select an existing volume (Disk1) and then right-click on the disk column. Select “Extend Volume” to proceed.

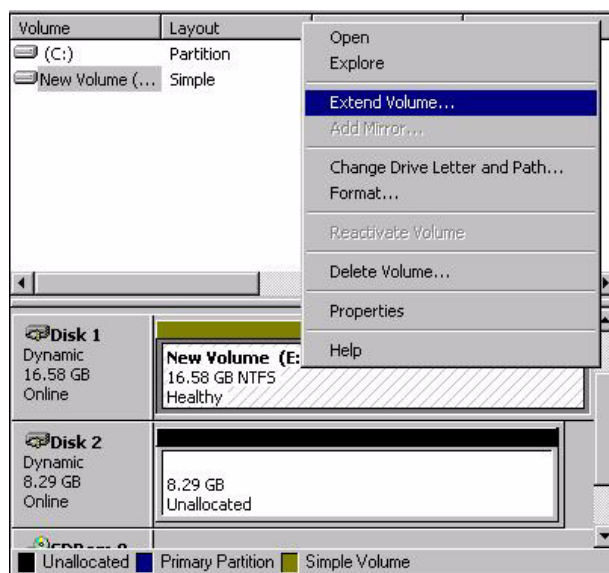


Fig. 11.28 Extend the Disk Volume on Windows

The Extend Volume Wizard will guide you through the rest of the process.

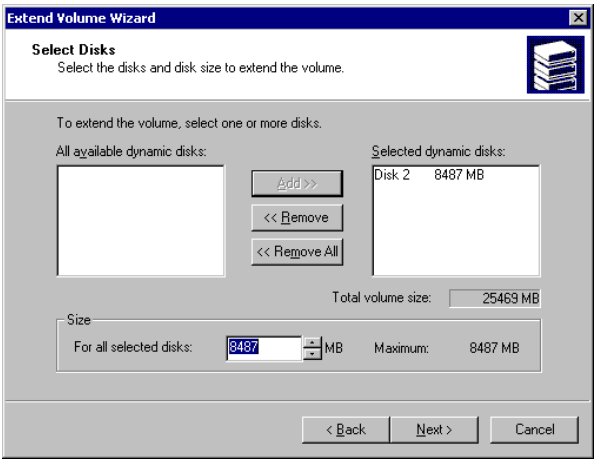


Fig. 11.29 Configure Disk Volume

The screen will display that the volume set of Drive E: has been extended into a spanned volume by the 8.3GB in Disk2.

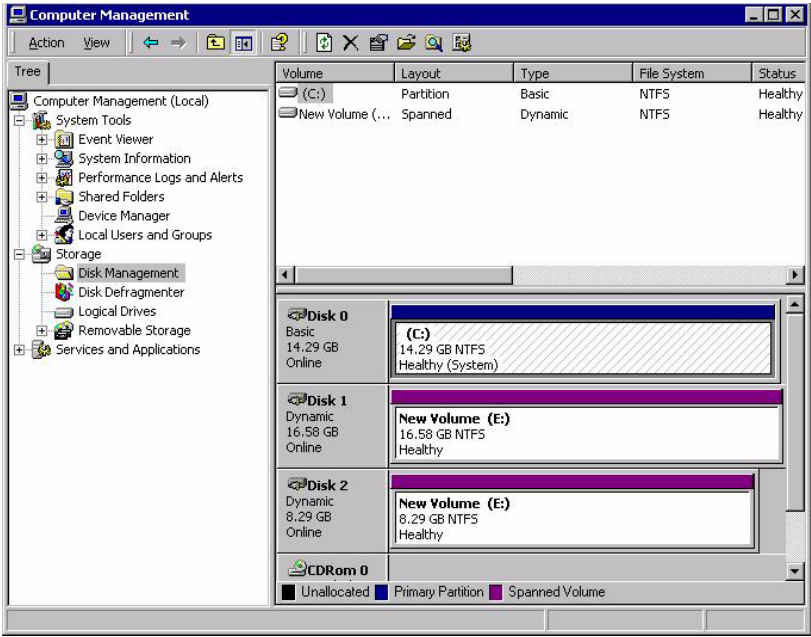


Fig. 11.30 View Disk Management on Windows

Logical Drive E: is now composed of two partitions with a total volume of 2500MB. To see this, hold down on the <Ctrl> key and select both Disk 1 and Disk2; then right-click your mouse and select “Properties.” Drive E: now has a capacity of about 25GB.

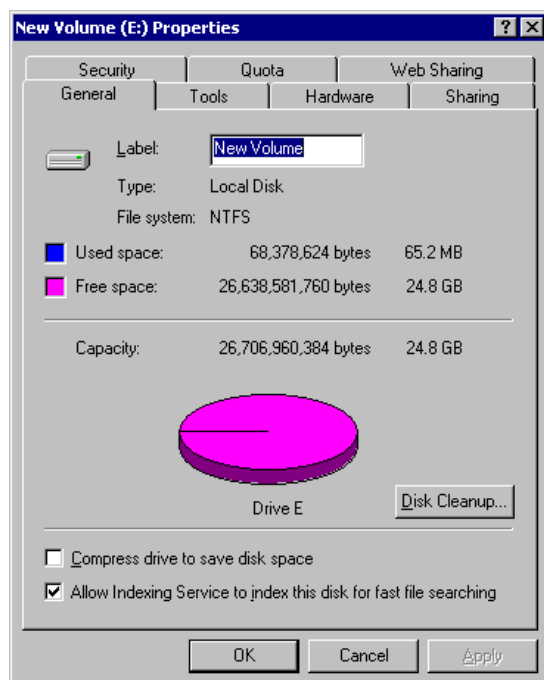


Fig. 11.31 View Disk Information

12S.M.A.R.T. Configuration

12.1Overview

S.M.A.R.T.

With the maturity of technologies like S.M.A.R.T., drive failures can be predicted to certain degree. Before S.M.A.R.T., receiving notifications of drive bad block reassignments may be the most common omen that a drive is about to fail. In addition to the S.M.A.R.T.-related functions as will be discussed later, a system administrator can also choose to manually perform “Clone Failing Drive” on a drive which is about to fail.

This function provides system administrators a choice on when and how to preserve data from a failing drive. Although not necessary under normal conditions, you may also replace any drive at-will even when the source drive is healthy.

The “Clone Failing Drive” can be performed under the following conditions:

1. Replacing a failing drive either detected by S.M.A.R.T. or notified by the controller.
2. Manually replacing and cloning any drive with a new drive.

12.2Clone Failing Drive

Unlike the similar functions combined with S.M.A.R.T., the “Clone Failing Drive” is a manual function. There are two options for cloning a failing drive: “Replace after Clone” and “Perpetual Clone.”

Replace after Clone

Data on the source drive, the drive with predicted errors (or any selected member drive), will be cloned to a standby spare and replaced later by the spare. The status of the replaced drive, the original member drive with predicted errors, will be redefined as a “used drive.” System administrators may replace the “used drive” with a new one, and then configure the new drive as a spare drive.

Locate the logical drive to which the drive with predictable errors belongs. Select the “Clone Failing Drive” function.

Select “Replace After Clone.” The controller will automatically start the cloning process using the existing “stand-by” (dedicated/global spare drive) to clone the source drive (the target member drive with predicted errors). If there is no standby drive (local/global spare drive), you need to add a new drive and configure it as a standby drive.

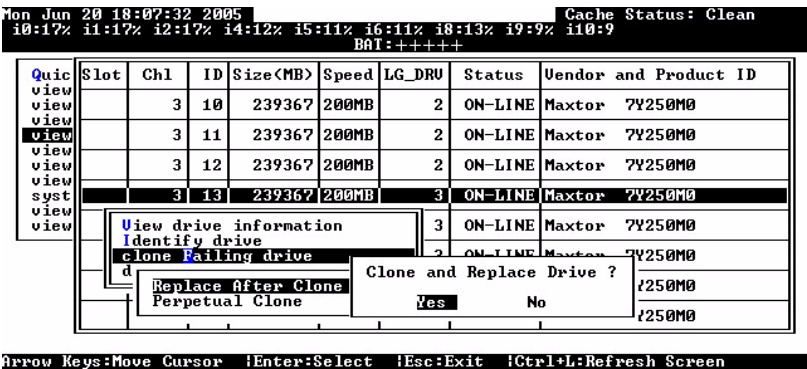


Fig. 12.1 Select Clone and Replace Drive




Fig. 12.2 Starting Clone and Replace Drive

The controller will automatically start the cloning process using the existing “stand-by” (local/global spare drive) to clone the source drive (the target member drive).
The cloning process will begin with a notification message:




Fig. 12.7 Notice Starting Clone Status

Press [ESC] to view the current progress:




Fig. 12.8 Cloning Progress Status

You may quit viewing the status bar by pressing [ESC] to return to the previous menu. Select the drive indicated as “CLONING” by pressing [ENTER]. Select “Clone Failing Drive” again to view the progress. You may identify the source drive and choose to “View Clone Progress” or “Abort Clone” if you happen to have selected the wrong drive.




Fig. 12.9 Identify The Source Drive

The cloning progress will be completed by a notification message as displayed below:




Fig. 12.10 Notice Clone Completed

You may press [ESC] to clear the notification message to see the drives’ status after the cloning process. The source drive (Channel 1 ID 5) remains as a member of logical drive “0,” and the “stand-by” drive (Channel 1 ID 2, the dedicated/global spare drive) has become a “CLONE” drive.




Fig. 12.11 View the Drive Status

12.3

S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology)

This section provides a brief introduction to S.M.A.R.T. as one way to predict drive failure and Bosch's implementations with S.M.A.R.T. for preventing data loss caused by drive failure.

Introduction

Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) is an emerging technology that provides near-term failure prediction for disk drives. When S.M.A.R.T. is enabled, the drive monitors predetermined disk drive attributes that are susceptible to degradation over time.

If a failure is likely to occur, S.M.A.R.T. makes a status report available so that the host can prompt the user to backup data from the failing drive. However, not all failures can be predicted. S.M.A.R.T. predictions are limited to the attributes the drive can monitor which are selected by the device manufacturer based on the attribute's ability to contribute to predict degrading or fault conditions.

Although attributes are drive specific, a variety of typical characteristics can be identified:

- Head flying height
- Data throughput performance
- Spin-up time
- Re-allocated sector count
- Seek error rate
- Seek time performance
- Spin try recount
- Drive calibration retry count

Drives with reliability prediction capability only indicate whether the drive is "good" or "failing." In a SCSI environment, the failure decision occurs on the disk drive and the host notifies the user for action. The SCSI specification provides a sense bit to be flagged if the disk drive determines that a reliability issue exists. The system then alerts the user/system administrator.

Bosch's Implementations with S.M.A.R.T.

Bosch uses the ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190 standard.

There are four selections related to the S.M.A.R.T. functions in firmware:

1. Disabled
Disables S.M.A.R.T.-related functions
2. Detect Only:
When the S.M.A.R.T. function is enabled, the controller will send a command to enable all drives' S.M.A.R.T. function, if a drive predicts a problem, the controller will report the problem in an event log.
3. Detect and Perpetual Clone
When the S.M.A.R.T. function is enabled, the controller will send a command to enable all drives' S.M.A.R.T. function. If a drive predicts a problem, the controller will report the problem in an event log. The controller will clone the drive if a Dedicated/Global spare is available. The drive with predicted errors will not be taken off-line, and the clone drive will still act as a standby drive.

If the drive with predicted errors fails, the clone drive will take over immediately. If the problematic drive is still working and another drive in the same logical drive fails, the clone drive will resume the role of a standby spare and start to rebuild the failed drive immediately. This prevents a fatal drive error if yet another drive should fail.

4. Detect and Clone + Replace

The controller will enable all drives' S.M.A.R.T. function. If a drive predicts a problem, the controller will report the problem in the form of an event log. The controller will then clone the problematic drive to a standby spare and take the problematic drive offline as soon as the cloning process is completed.

**NOTICE!**

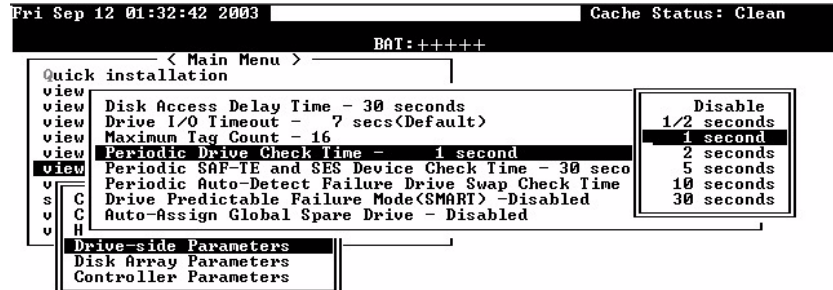
If you are using drives of different brands in your RAID system, as long as they are ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190-compatible, there should not be any problems working with the controller/subsystem.

12.4 Configuration Procedure

Enabling the S.M.A.R.T. Feature

Follow the procedure below to enable S.M.A.R.T. on all drives.

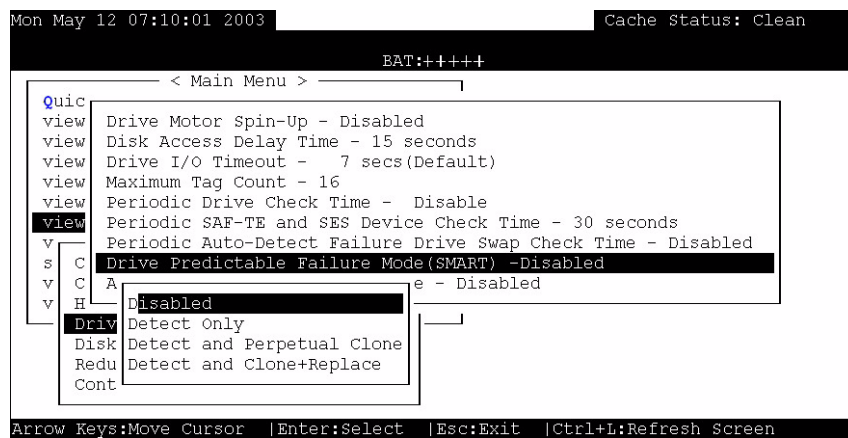
1. First, enable the "Periodic Drive Check Time" function. In \View and Edit Configuration Parameters\Drive-side Parameters\Periodic Drive Check Time, choose a time interval.



Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 12.12 Select Drive Check Time

2. In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose one from "Detect Only," "Detect, Perpetual Clone" and "Detect, Clone+Replace."



Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 12.13 Select Drive Failure Mode

Examining Whether Your Drives Support S.M.A.R.T.

To see if your drive supports S.M.A.R.T., follow the steps below:

1. Enable "S.M.A.R.T." for your drives in the RAID system.
2. In "View and Edit Drives," choose one drive to test to. Press [ENTER] on the drive; a sub-menu will appear.
3. Note that a new item "Predictable Failure Test" appears in the sub-menu. If the S.M.A.R.T. feature is not properly enabled, this item will not appear in the sub-menu.

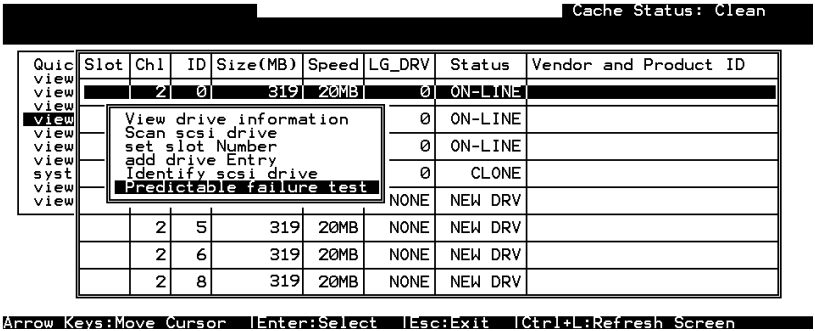


Fig. 12.14 Predictable Failure Test

- Choose “Predictable Failure Test,.” The controller will force the drive to simulate predictable drive errors.





Fig. 12.15 Choose Predictable Failure Testing

- Press [ENTER], and after a while (the next time the controller performs “Periodic Drive Check”), the controller will detect the errors simulated by the drive. An error message displays like this: “SMART-CH:? ID:? Predictable Failure Detected (TEST).” If this error message appears, it means your drive supports S.M.A.R.T. features.



- If the error message does not appear, you may simply refer to related documentation or contact your drive manufacturer for information about whether the drive model and drive firmware version support S.M.A.R.T.

Using S.M.A.R.T. Functions

- Enable “S.M.A.R.T.” on the RAID controller.
- Make sure your drives do support S.M.A.R.T. so that your system will work properly.
- The “Detect Only” Setting
In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose “Detect Only.”

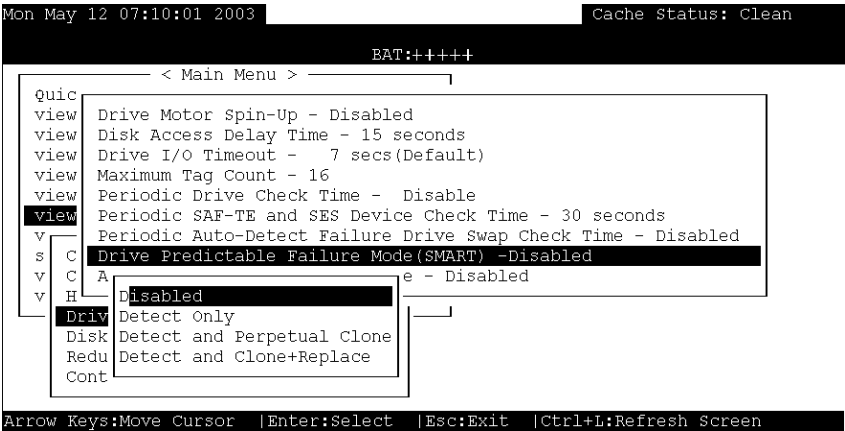


Fig. 12.16 Choose Detect Only on SMART Function

Whenever a drive detects symptoms of predictable drive failure, the controller will issue an error message.

4. The “Detect, Perpetual Clone” Setting

In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose “Detect, Perpetual Clone.”

You should have at least one spare drive for the logical drive (either Local Spare or Global Spare Drive).

When a drive (logical drive member) detects predictable drive errors, the controller will “clone” the drive with a spare drive. You may enter the "View and Edit Drives" menu and click on the spare drive (either Local or Global one). Choose from the menu items if you want to know the status of the source drive, the cloning progress, or to abort cloning.

Quick view

Fig. 12.17 Choose Source Drive Status



NOTICE!

As a precaution against the untimely failure of yet another drive, when configured as “perpetual clone,” the spare drive will only stay mirrored to the source drive (the drive with signs of failure), but not replace it until the source drive actually fails.

While the spare drive is mirroring the source drive, any occurrence of drive failure (when there are no other spare drives) will force the spare drive to give up the mirrored data and resume its original role – it will become a spare drive again and start rebuilding the failed drive.

5. The “Detect, Clone + Replace” Function

In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose “Detect, Clone+Replace.”

Make sure you have at least one spare drive to the logical drive. (Either Local Spare Drive or Global Spare Drive)

When a drive (a logical drive member) detects the predictable drive failure, the controller will “clone” the drive with a spare drive. After the “clone” process is completed, it will replace the source drive immediately. The source drive will be identified as a “used drive.”

If you want to see the progress of cloning, press [ESC] to clear the notification message and see the status bar.

The source drive's status will be defined as a "used drive" and will be immediately replaced and pulled offline. This drive should be replaced with a new one as soon as possible.

13

Implementations for AV Applications

This chapter introduces some new firmware functions for AV applications. More options will be available for AV applications with future firmware releases.



NOTICE!
All figures in this chapter are showing examples using the management hyper terminal screen.

13.1

Maximum Drive Response Time

In situations such as drive failure or the occurrence of media error, a read or write action returned after several hundreds milliseconds would prove too long and cause choppy audio or dropped video frames.

The maximum response time option, as displayed below, provides a timeout value for processing read/write requests. If timeout is reported on a specific member of an array, the subsystem immediately retrieves data from the parity data and other members of the array. In this way, causes of delay can be eliminated.

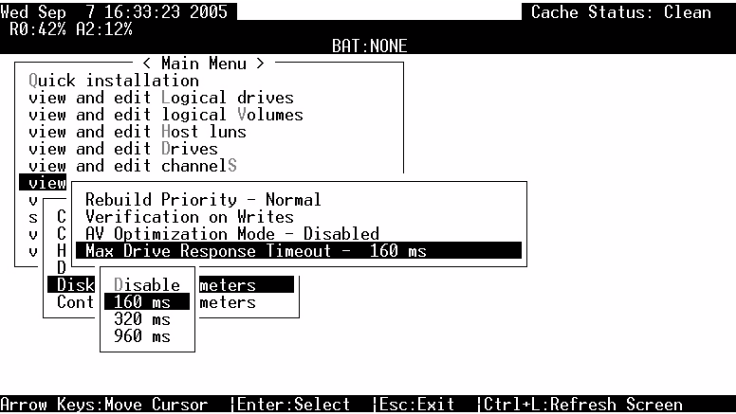


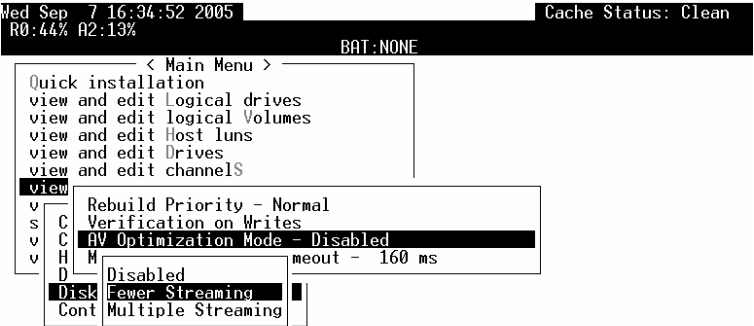
Fig. 13.1 Selecting Maximum Drive Response Timeout Value

- To prepare the array for read-intensive applications, the following are recommended:
1. Performance using the write-through caching mode is better than that using the write-back mode on subsystems equipped with redundant RAID controllers.
 2. Arrays should not be partitioned.
 3. The priorities for Rebuild or Add-drive operations should be set to “low.”
 4. Another timeout value, the “Drive I/O Timeout” which determines whether a drive has eventually failed to respond to I/O requests, is required as the first-level timeout.

13.2

AV Optimization Mode

The AV optimization is applied for the emerging streaming Audio/Video applications. This option prepares the array for applications requiring continuous reading such as video streaming.



Arrow Keys:Move Cursor |Enter:Select |Esc:Exit |Ctrl+L:Refresh Screen

Fig. 13.2 Selecting An AV Optimization Mode

The applicable number of streaming for fewer/multiple streaming will depend on the field applications.

i

NOTICE!

No matter what user configures for the Maximum Drive Response Timeout, once the AV optimization is set, the Maximum Drive Response Timeout is fixed at its predefined value.

14 Firmware Functionality Specifications

14.1 Basic RAID Management

Specification	Feature
RAID levels	0, 1(0+1), 3, 5, 10, 30, 50, JBOD and NRAID. Levels 10, 30 and 50 are the multi-level RAID with the logical volume implementations; logical volumes consist of logical drives of different RAID levels that are striped together.
Maximum number of logical drives	16 as default, up to 32 with a 1GB or above memory size
Maximum logical drive capacity	64TB – when optimization mode is set to Sequential 16TB – when optimization mode is set to Random
RAID level dependency to each logical drive	Independent. Logical drives configured in different RAID levels can co-exist in a logical volume and within a RAID subsystem
Maximum number of logical drive members	128 - 512MB memory size 112 - 256MB memory size
Configurable stripe size	4KB, 16KB, 32KB, 64KB, 128KB, 256KB, 512KB, or 1024KB per logical drive <ul style="list-style-type: none"> 128KB is the default for earlier firmware with optimization for Sequential I/Os 32KB is the default for earlier firmware with optimization for Random I/Os
Configurable Write Policy (write policy per array)	Write-Back or Write-Through per logical drive. This policy can be modified later.
Logical drive identification	Unique, controller randomly generated logical drive ID; Logical drive name user-configurable for ease of identification in a multiple arrays configuration
Maximum number of partitions for each logical drive	16 as default, up to 64 with a 1GB memory size
Maximum number of logical volumes	8 as default, up to 16 with a 1GB or above memory size
Maximum number of LUNs Mappable	Default is 128, up to 1024 with a 1GB or above memory size
Maximum number of LUNs per host ID	Up to 32, user configurable
Maximum number of Media Scan task schedules	16
Concurrent I/O	Supported
Tag Command Queuing (TCQ)	Supported
Native Command Queuing (NCQ)	Supported
Dedicated spare drive	Supported, hereby defined as the spare drive specifically assigned to a logical drive

Specification	Feature
Global spare drive	Supported, the spare drive that serves all logical drives (as long as it is equal in size or larger than logical drive members)
Global spare auto-assign	Supported, applies to non-configured drive(s); safeguards the array if a spare has been used in the previous array rebuild and users forget to configure a new drive as a spare.
Co-existing dedicated and global spare drives	Supported
Auto-rebuild onto spare drive	Supported
Auto-scan of replacement drive upon manually initiated rebuild	Supported
One-step rebuild onto a replacement drive	Supported
Immediate logical drive availability	Supported; Logical arrays are immediately ready for Host I/Os. Initialization task is completed in the background except when the logical array is stated as "INCOMPLETE" or "BAD;" e.g., has a failed member right after the creation.
Auto-rebuild onto failed drive replacement	Supported. With no spare drive, the subsystem will auto-scan the failed drive and starts rebuild automatically once the failed drive has been replaced.
Concurrent rebuild / expansion	Multiple logical drives can proceed with a Rebuild/Regenerating Parity, and/or Expansion/Initialization/Add Drive operation at the same time. NOTE: Regenerate Parity and Rebuild cannot take place on a logical drive at the same time. Create, Expand, and Add Drive operations cannot take place on a logical drive at the same time.
Background firmware download	Firmware can be downloaded during active I/Os.
Auto recovery from logical drive failure (configuration on drives)	Supported. When user accidentally removed the wrong drive to cause the 2nd drive failure of a one-drive-failed RAID5 / RAID3 logical drive, fatal error may occur. However, you may force the subsystem to reaccept the logical drive by switching off the subsystem, installing the drive back to its original drive slot, and then power on the subsystem. The logical drive will be restored to the one-drive-failed status.

NOTICE!

To ensure smooth operation, sufficient cache memory buffer is required for configurations made up of numerous logical units. An intelligent trigger mechanism is implemented with firmware version 3.47 and later. If a subsystem/controller comes with a DIMM module of the size equal or larger than 1GB, firmware automatically enlarges the maximum numbers of logical units.

	DIMM size < 1GB	DIMM size >= 1GB
Max. no. of LD	16	32
Max. no. of LV	8	16
Max. partitions per LD/LV	16	64
Max. no. of LUN	128	1024

14.2 Advanced Features

Media Scan	Supported. Verify written data on drives to avoid bad blocks from causing data inconsistency. If bad blocks are found, data can be reconstructed by comparing and recalculating parity from adjacent drives (RAID1/3/5). The "Reconstruction Writes" are followed by "Write Verification" operation.
Bad Block Handling in degraded mode	A method for handling low quality drives. The operation is performed on both the logical drive in degraded mode or that being rebuilt. If bad blocks should be encountered during Rebuild, Add Drive, Host Write, or Regenerate Parity operation, the controller will first attempt to reconstruct affected data and those unrecoverable bad blocks are stated as bad and passed to host. Users have the option to abandon data on the unrecoverable sectors to continue rebuild in a degraded mode. Low quality drive handling comes with transparent resetting of hung hard drives.
Transparent reset of hung HDDs	Supported
Auto cache flush on critical conditions (caching mode dynamic switch)	When critical conditions occur, e.g., component failure, or BBU under charge, cached data will be flushed and the write policy will be changed to write-through mode. Configurable "Trigger Events" for Write-through/Write-Back Dynamic Switch. The configuration can also be set with the Oem "lappend" utility.
Drive low-level format	Supported
RAID parity update tracking and recovery	Yes
Host-side Ordered Tag support	Supports write commands with embedded Ordered Tags
Drive identification (flash drive function)	Supported. Force a drive to light on its activity indicator for users to recognize its location
Drive information listing	Supported. Drive vendor name, model number, firmware revision, capacity (blocks), serial number, narrow/wide and current sync. speed
Drive read/write test	Supported
Configuration on disks	Supported. The logical drive information is recorded on drive media. The logical drives can still be accessed if using different Bosch RAID controllers/subsystems, e.g., drives removed and installed in a different subsystem.
Save/ restore NVRAM to / from disks	Supported. Save all the settings stored in the controller NVRAM to the logical drive members
Save / restore NVRAM to / from a file	Supported. Save all the settings stored in the controller NVRAM to a file (via GUI manager) on user's computer
Host-side 64-bit LBA support	Supports array configuration (logical drive, logical volume, or a partition of them) of a capacity up to 64TB.

Host LUN geometry: user configurable default geometry (Solaris OSes)	<ul style="list-style-type: none"> Capacity <64GB: Head=63, Sector=32, Cylinder=? (depends on capacity) 64GB<capacity<128GB: Head=64, Sector=64, Cylinder=? (depends on capacity) 128GB<capacity<256GB: Head=127, Sector=64, Cylinder=? (depends on capacity) 256GB<capacity<512GB: Head=127, Sector=127, Cylinder=? (depends on capacity) 512GB<capacity<1TB: Head=255, Sector=64, Cylinder=? (depends on capacity) 1TB<capacity: Head=225, Sector=225, Cylinder=? (depends on capacity) For capacity larger than 1TB, please refer to <i>Section 9 Enclosure Management</i>
User configurable geometry range:	Sector: 32, 64, 127, 255 or Variable Head: 64, 127, 255 or Variable Cylinder: <1024, <32784, <65536 or Variable
Drive motor spin-up	Supported. The controller will send spin-up (start unit) command to each drive at the 4 sec. intervals.
Drive-side tagged command queuing	Supported. User adjustable up to 128 for each drive
Host-side maximum queued I/O count	User adjustable up to 1024
Maximum concurrent host LUN connection	User adjustable up to 64
Number of tags reserved for each Host-LUN connection	User adjustable up to 256
Controller shutdown	Flushes cached contents upon the detection of critical conditions, e.g., a high temperature condition persists for a long time
Drive I/O timeout	User adjustable
I/O channel diagnostics	Supported; please contact your dealer for more details
Maximum Drive Response Time (Guaranteed Latency I/O)	User adjustable from 160 to 960ms. If a disk drive fails to return data on read requests before the timeout value is exceeded, the array immediately generates data from the parity data and the other members of a logical drive.
Drive roaming	Supported; array configuration stored on disk drives

14.3 Caching Operation

Write-back cache	Supported.		
Write-through cache	Supported.		
Supported memory type	SDRAM or DDR memory for enhanced performance Fast Page Memory with Parity for enhanced data security		
Read-ahead operation	Intelligent and dynamic read-ahead operation for processing sequential data requests		
Multi-threaded operation	Yes		
Scatter / Gather	Supported		
I/O sorting	Supported. Optimized I/O sorting for enhanced performance.		
Adaptive Write-back/Write-through switching	For a better performance when handling large sequential writes, firmware temporarily disables write-cache and the synchronized cache operation between partner controllers if operating with dual-active RAID controllers. Firmware automatically restores the write-back mode when later encountering random and small writes.		
Periodic Cache Flush	<p>Firmware can be configured to flush the cached contents in memory at every preset interval:</p> <ol style="list-style-type: none"> 1. If data integrity is of the concern, e.g., the lack of a battery backup protection. 2. Cache flush on preset intervals to avoid the latency when cache memory is full due to write delays. 		
Variable stripe size		Opt. for Sequential I/Os	Opt. for Random I/Os
	RAID0	128	32
	RAID1	128	32
	RAID3	16	4
	RAID5	128	32
Caching Optimization	<ul style="list-style-type: none"> • Cache buffer sorting prior to cache flush operation • Gathering of writes during flush operation to minimize the number of I/Os required for parity update • Elevator sorting and gathering of drive I/Os • Multiple concurrent drive I/Os (tagged commands) • Intelligent, predictive multi-threaded read-aheads • Multiple, concurrent host I/O threads (host command queuing) 		

14.4 RAID Expansion

On-line RAID expansion	Supported. Capacity brought by array expansion is immediately ready for Host I/Os when its status changes from "EXPAND" to "INITIALIZING." Initialization task is then completed in the background except when the logical array is stated as "INCOMPLETE" or "BAD;" e.g., has a failed member right after creation.
Mode-1 RAID expansion -add drive	Supported. Multiple drives can be added concurrently. Add Drive can even be performed in degraded mode.
Mode-2 RAID expansion – copy and replace drives	Supported. Replace members with drives of larger capacity.
Expand capacity with no extra drive bays required	Supported in Mode 2 RAID expansion, which provides "Copy and Replace Drive" function to replace drives with drives of greater capacity. No need for hardware upgrade, e.g., adding a new enclosure for the extra drives.
Operating system support for RAID expansion	No. No operating system driver required. No software need to be installed for this purpose.

14.5 S.M.A.R.T. Support

Copy & replace drive	Supported. User can choose to clone a member drive showing symptoms of defects before it fails.
Drive S.M.A.R.T. support	Supported, with intelligent error handling implementations.
User selectable modes for S.M.A.R.T.	<ul style="list-style-type: none"> • Detect only • Perpetual Clone: using a hot-spare to clone the drive reporting SMART errors; the hot-spare remains a clone drive • Clone + Replace: using a hot-spare to replace the drive reporting SMART errors; the drive reporting errors is pulled offline

14.6 Redundant Controller

Active-active redundant controller	Supported
Synchronized cache	Supported. Through one or multiple, dedicated synchronizing channels on a common backplane or external cabling. Synchronized cache over SCSI channels, Fibre loops, or SATA channels is supported. Synchronized cache can be disabled when using write-through mode in a redundant controller configuration to prevent performance trade-offs.
Write-back cache enabled in redundant controller mode	Yes, with synchronized cache connection between controllers.
Automatic failover	Yes (user's interaction necessary)
Automatic failback	Yes (user's interaction necessary)
Controller hot-swap	<ul style="list-style-type: none"> No need to shut down the failed controller before replacing the failed controller. Support online hot-swap of the failed controller. There is no need to reset or shutdown the failed controller. One controller can be pulled out during active I/Os to simulate the destructive controller failure. (Customer's design-in hot-swap mechanism necessary)
Parity synchronization in redundant controller write-back mode to avoid write-hole	Supported.
Redundant controller communication over Fibre loops	Dedicated loops or distribution over drive loops selectable
No single-point-of-failure	Supported.
Automatic engagement of replacement controller	Supported.
Dynamic cache memory allocation	Yes. Cache memory is dynamically allocated, not fixed.
Environment management	Supported. SAF-TE, S.E.S., and on-board controller voltage/temp monitor are all supported in both single and redundant controller mode. In the event of controller failure, services can be taken over by the surviving controller.
Cache battery backup	Supported. Battery backup modules support cached data during a power outage in both single controller and redundant controller modes.
Load sharing	Supported. Workload can be flexibly divided between different controllers by assigning logical configurations of drives (LDs/LVs) to different controllers.
Require a special firmware for redundant controller?	No.

14.7

Data Safety

Regenerate parity of logical drives	Supported. Can be manually executed to ensure that bad sectors do not cause data loss in the event of drive failure.
Scheduled Media Scan	Revision 3.34 and above allows Media Scan to be scheduled starting at a specified start time and repeated at regularly timed intervals. The start time and time intervals can be selected from drop-down menus. Start time is manually entered using its numeric representatives in the following order [MMDDhhmm[YYYY]], and it reads the date and time set for the controller's real-time clock. The selectable time intervals (the Execution Period) range from one (1) second to seven (7) weeks. Each such schedule can be defined to operate on individual hard drives, all members of a specified logical drive, or members of selected logical drives. Each schedule can include up to five (5) logical drives. The RS-232C terminal interface and RAIDWatch revision 2.0 support this functionality.
Bad block auto-reassignment	Supported. Automatic reassignment of bad block
Battery backup for cache memory	Supported. The battery backup solutions provide long-lasting battery support to the cache memory when power failure occurs. The unwritten data in the cache memory can be committed to drive media when power is restored.
Verification on normal writes	Supported. Performs read-after-write during normal write processes to ensure data is properly written to drives.
Verification on rebuild writes	Supported. Performs read-after-write during rebuild write to ensure data is properly written to drives.
Verification on LD initialization writes	Supported. Performs read-after-write during logical drive initialization to ensure data is properly written to drives.
Drive S.M.A.R.T. support	Supported. Drive failure is predictable with reference to the variables detected. Reaction schemes are selectable from Detect only, Perpetual Clone and Copy + Replace. These options help to improve MTBF.
Clone failing drive	Users may choose to clone data from a failing drive to a backup drive manually.
Automatic shutdown on over-temperature condition	Controller automatically starts a shutdown sequence (entering an idle state receiving no I/O requests) upon the detection of high-ambient temperature for an extended period of time.

14.8 System Security

Password protection	Supported. All settings require the correct password (if set) to ensure system security.
User-configurable password validation timeout	Supported. After certain time in absence of user interaction, the password will be requested again. This helps to avoid unauthorized operation when user is away.
SSL-enabled RAIDWatch Agents	Agents communicate to the controller through limited set of authorization options.

14.9 Environment Management

SAF-TE/S.E.S. support	Supported. The SAF-TE/S.E.S. modules can be connected to the drive channels. The RAID controller will detect errors from SAF-TE/S.E.S. devices or notify drive failures via SAF-TE/S.E.S. <ul style="list-style-type: none"> Both SAF-TE/S.E.S. via drive and device-self-interfaced methods are supported. Redundant SAF-TE/S.E.S. devices are supported Multiple S.E.S. devices are supported
Dynamic on-lining of enclosure services	Once an expansion unit (JBOD) with supported monitoring interface is combined with a RAID system, its status will be automatically polled.
SAF-TE/S.E.S. polling period	User configurable (50ms, 100ms, 200ms, 500ms, 1~60sec)
Multiple SAF-TE/S.E.S. modules on the same channel	Supported.
Multiple SAF-TE /S.E.S. modules on different channels	Supported.
Mapping SAF-TE/S.E.S. device to host channel for use with host-based SAF-TE/S.E.S. monitoring	Supported.
Event Triggered Operation	When any of the following happens, the firmware disables write-back caching to minimize the chance of losing data: <ul style="list-style-type: none"> Battery, controller, cooling fan, or PSU failure The upper temperature thresholds are exceeded Low battery charge The triggering factors are user-configurable
Dual-speed cooling fan control	Yes
Dual-LED drive status indicators	Supported. Both single-LED and dual-LED drive status indicators are supported.
SAF-TE/ S.E.S. temperature value display	Supported. Display the temperature value provided by enclosure SAF-TE module (if available).
On-board controller voltage monitors	Supported. Monitors the 3.3V, 5V, and 12V voltage status. Event triggered thresholds user configurable.

On-board controller temperature sensors	Supported. Monitors the CPU and board temperature status. Event trigger threshold user configurable.
Enclosure redundant power supply status monitoring	Supported. SAF-TE/S.E.S.
Enclosure fan status monitoring	Supported. SAF-TE/S.E.S
Enclosure temperature monitoring	Supported. SAF-TE/S.E.S

14.10 User Interface

RAIDWatch on-board	Out-of-band configuration and monitoring via Ethernet. Browser accessible configuration options by installing RAID-Watch to reserved space on drive via ftp.
RS-232C terminal	Supports terminal modes: ANSI, VT-100, ANSI Color. Provides menu-driven user-friendly text-based interface.
Graphical user interface (Java-based GUI manager)	Provides user-friendly graphical interface. Communicates with RAID controller via Out-of-band Ethernet, In-band SCSI, In-band Fibre or SNMP (Windows-based GUI).
SSH support	Secure Shell over Telnet supported
External interface API for customized host-based management	Supported.
LCD front panel	Provides easy access for user instinct operation.
Buzzer alarm	Warns users when any failures or critical events occur.

14.11 High Availability

Custom inquiry serial number	Custom Inquiry Serial Number (for support of multi-pathing software like Veritas, QLogic, etc)
-------------------------------------	--

15 System Functions - Upgrading Firmware

15.1 Upgrading Firmware

The RAID controller's firmware resides in flash memory that can be updated through the COM port, LAN port, or via In-band SCSI/Fibre.

15.1.1 Background RS-232C Firmware Download: Single Controller

Host I/Os will not be interrupted during the download process. After the download process is completed, user should find a chance to reset the controller for the new firmware to take effect.

15.1.2 Note for Redundant Controller Firmware Upgrade:

A controller used to replace a failed unit in a dual-controller system is often running a newer release of firmware version. To solve the contention, make sure the firmware on a replacement controller is downgraded to that running on the surviving controller.

NOTICE!

- Allow the downloading process to finish. Do not reset or turn off the computer or the controller while it is downloading the file. Doing so may result in an unrecoverable error that requires the service of the manufacturer.
- When upgrading the firmware, check the boot record version that comes with it. If the boot record version is different from the one installed on the surviving controller previously, the new boot record binary must be installed.
- The COM 2 serial port cannot be used to download firmware.

15.2 Upgrading Firmware Using RAIDWatch Manager

15.2.1 Establish the In-band SCSI connection in RAIDWatch Manager

Please refer to RAIDWatch User's Manual for details on establishing the management session for RAIDWatch Manager.

15.2.2 Upgrade Both Boot Record and Firmware Binaries

1. Connect to the RAID system locally or from a remote computer using RAIDWatch Manager. While connected to the RAID system, there will be icon(s) with IP address specified on the left of the menu screen. Select by double-clicking the icon of the RAID system which firmware is to be upgraded. Select the "Configuration Parameters" icon and then select the "System" tab. Single-click the "Download FW+BR" check circle. Click the Apply button to proceed. A message prompt should display. Click Yes and then a file location prompt will appear.
2. Provide the boot record binary filename, the RAIDWatch Manager will start to download the boot record binary to the controller.
3. After the boot record download is completed, provide the firmware filename to the RAIDWatch Manager. It will start to download the firmware to RAID controller.
4. Shutdown the application server(s) or stop the application(s) which are accessing the RAID, then reset the RAID subsystem/controller in order to use the new downloaded firmware. With firmware release 3.21 and above, host I/Os will not be interrupted by the download process. Users may find a chance to stop host I/Os and reset the controller for new firmware to take effect.

15.2.3 Upgrade the Firmware Binary Only

1. Connect to the RAID system locally or from a remote computer using RAIDWatch Manager. While connected to the RAID system, there will be icon(s) with IP address specified on the left of the menu screen. Select by double-clicking the icon of the RAID subsystem which firmware is to be upgraded. Select the "Configuration Parameters" icon and then select the "System" tab. Single-click the "Download FW" check circle. Click the Apply button to proceed. A message prompt should display. Click Yes and then a file location prompt will appear.
2. Provide the firmware filename to the RAIDWatch Manager. It will start to download the firmware to RAID controller.
3. Shutdown the application server(s) or application(s) which are accessing the RAID, then reset the controller in order to use the new downloaded firmware.

15.3 Upgrading Firmware Using RS-232C Terminal Emulation

The firmware can be downloaded to the RAID controller/subsystem by using an ANSI/VT-100 compatible terminal emulation program. Whichever terminal emulation program is used must support the ZMODEM file transfer protocol. The following example uses the HyperTerminal in Windows NT®. Other terminal emulation programs (e.g., Telix and PROCOMM Plus) can perform the firmware upgrade as well.

15.3.1 Establishing the connection for the RS-232C Terminal Emulation

Please refer to *Section 3 Accessing the Array: Serial Port, Ethernet, and Access Security*.

15.3.2 Upgrading Both Boot Record and Firmware Binaries

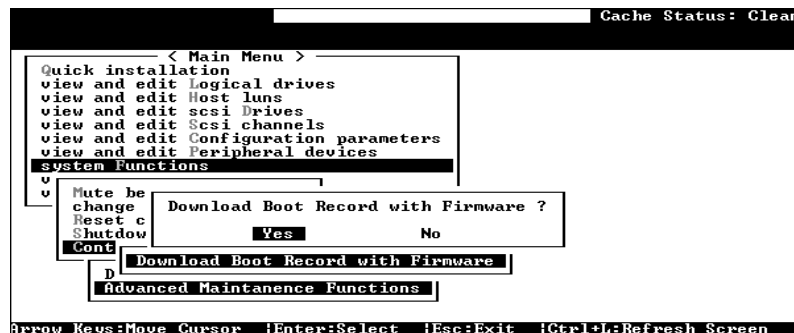


Fig. 15.1 Upgrading Boot Record and Firmware Binaries Using Hyper Terminal Screen

1. From the Main Menu, scroll down to "System Functions."
2. Go to "Controller Maintenance."
3. Choose "Advanced Maintenance."
4. Select "Download Boot Record and Firmware."
5. Set ZMODEM as the file transfer protocol of your terminal emulation software.
6. Send the Boot Record Binary to the controller. In HyperTerminal, go to the "Transfer" menu and choose "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).
7. After the Boot Record has been downloaded, send the Firmware Binary to the controller. In HyperTerminal, go to the "Transfer" menu and choose "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).
8. When the Firmware completes downloading, the controller will automatically reset itself. For a newer version of firmware, you need to manually reset the subsystem/controller for the new firmware to take effect.

15.3.3

Upgrading the Firmware Binary Only

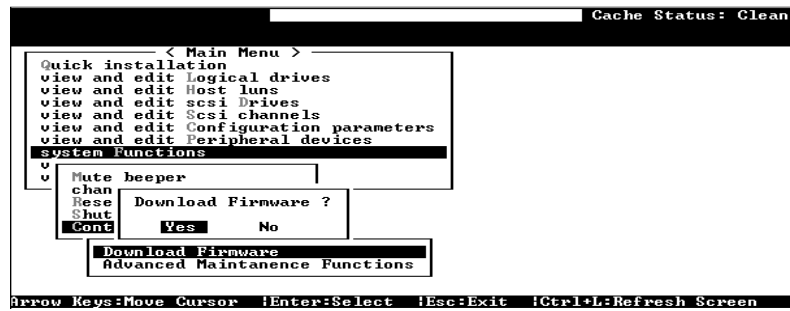


Fig. 15.2 Upgrading Firmware Binaries Using Hyper Terminal Screen

1. From the Main Menu, scroll down to "System Functions."
2. Go to "Controller Maintenance."
3. Choose "Download Firmware."
4. Set ZMODEM as the file transfer protocol of your terminal emulation software.
5. Send the Firmware Binary to the controller. In Hyper Terminal, select "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).

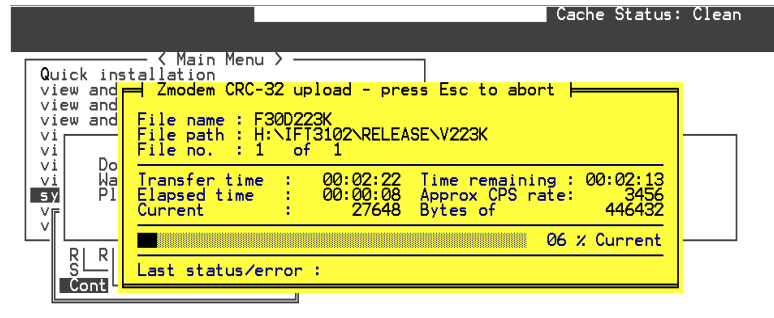


Fig. 15.3 Sending Files

6. When the Firmware completes downloading, the controller will automatically reset itself. If the subsystem/controller is running firmware later than version 3.21, you need to manually reset the subsystem/controller for the new firmware to take effect.

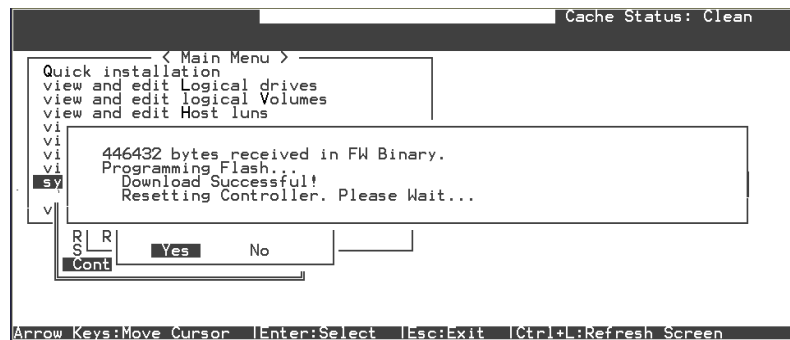


Fig. 15.4 Receiving Firmware Files

16 Event Messages

The controller events can be categorized as follows according to the severity levels:

Critical	Errors that need to attend to immediately
Warning	Errors
Notification	Command processed message sent from Firmware

The RAID subsystem records all system events from power on, it can record up to 1,000 events. To power off or to reset the controller will cause an automatic deletion of all the recorded event logs.

The RAIDWatch manager can be used to record events on multiple subsystems especially when controller reset or power-off is an expected action. The events can also be seen from RAIDWatch's Configuration Client utility. Associated details can be found in the RAIDWatch user's manual and online help.

Descriptions below may contain abbreviations. Abbreviations and Capitalized letters are preserved for the coherency with the event messages shown on LCD screen or terminal.

16.1 Event Index

16.1.1 Logical Drive Events

Critical

LG:0 Logical Drive ALERT: CHL:0 ID:132 Drive Missing
 LG:0 Logical Drive ALERT: CHL:0 ID:132 Drive Failure
 LG:0 Logical Drive ALERT: Rebuild Aborted
 LG:0 Logical Drive ALERT: Rebuild Failed
 LG:0 Logical Drive ALERT: Parity Regeneration Aborted
 LG:0 Logical Drive ALERT: Parity Regeneration Failed
 LG:0 Logical Drive ALERT: Creation Aborted
 LG:0 Logical Drive ALERT: Creation Failed
 LG:0 Logical Drive ALERT: Initialization Failed
 LG:0 Logical Drive ALERT: Expansion Aborted
 LG:0 Logical Drive ALERT: Expansion Failed
 ALERT: CHL:0 ID:132 Media Scan Failed
 LG:0 ALERT: CHL:0 ID:132 Media Scan Failed
 LG:0 Logical Drive ALERT:CHL:0 ID:132 Clone Failed
 LG:0 Logical Drive ALERT:CHL:0 ID:132 Clone Aborted
 LG:0 Logical Drive ALERT:CHL:0 ID:132 Clone Failed
 LG:0 Logical Drive ALERT:CHL:0 ID:132 Clone Aborted
 LG:0 Logical Drive ALERT: Logical Drive Bad Block Table FULL
 LG:0 Logical Drive ALERT: Logical Drive Bad Block Table BAD
 LG:0 Logical Drive ALERT: Logical Drive On-Line Init Table BAD
 LG:0 Logical Drive ALERT: Bad Data Block Marked 000000084
 LG:0 Logical Drive ALERT: UNPROTECTED Block Marked 000000084
 LG:0 Logical Drive NOTIFY: Bad Data Block Encountered 000000084
 LG:0 ALERT: Inconsistent Parity Encountered Block 20000000B

Notification

LG:0 Logical Drive NOTICE: Continue Rebuild Operation
 LG:0 Logical Drive NOTICE: Starting Rebuild
 LG:0 Logical Drive NOTICE: Rebuild of Logical Drive Completed

LG:0 Logical Drive NOTICE: Continue Parity Regeneration Operation
LG:0 Logical Drive NOTICE: Starting Parity Regeneration
LG:0 Logical Drive NOTICE: Parity Regeneration of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting Creation
LG:0 Logical Drive NOTICE: Creation of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting On-Line Initialization
LG:0 Logical Drive NOTICE: On-Line Initialization of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting Off-Line Initialization
LG:0 Logical Drive NOTICE: Off-Line Initialization of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting On-Line Expansion
LG:0 Logical Drive NOTICE: On-Line Expansion of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting Off-Line Expansion
LG:0 Logical Drive NOTICE: Off-Line Expansion of Logical Drive Completed
LG:0 Logical Drive NOTICE: Continue Add Drive Operation
LG:0 Logical Drive NOTICE: Starting Add Drive Operation
LG:0 Logical Drive NOTICE: Add Drive Operation Paused
LG:0 Logical Drive NOTICE: Add Drive to Logical Drive Completed
LG:0 Logical Drive NOTICE: Add Drive Operation Paused
NOTICE: CHL:0 ID:132 Starting Media Scan
LG:0 NOTICE: CHL:0 ID:132 Starting Media Scan
NOTICE: Media Scan of CHL:0 ID:132 Completed
LG:0 NOTICE: Media Scan of CHL:0 ID:132 Completed
LG:0 ALERT: CHL:0 ID:132 Media Scan Aborted
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Continue Clone Operation
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Starting Clone
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Copy and Replace Completed
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Clone Completed
LG:0 Logical Drive NOTIFY: Bad Data Block Recovered 000000084
LG:0 Logical Drive NOTIFY: Bad Data Block Recovered 000000084

16.1.2 Channel and Individual Drive Events

Critical

CHL:0 ID:132 ALERT: Media Error Unrecoverable-0xD8001C7C
CHL:0 FATAL ERROR(0)

Warning

CHL:1 ID:0 Target ALERT: Unexpected Select Timeout
CHL:1 ID:0 Target ALERT: Unexpected Select Timeout
CHL:1 Drive Channel ALERT: Unexpected Select Timeout
CHL:1 RCC Channel ALERT: Gross Phase/Signal Error Detected
CHL:1 ID:0 Target ALERT: Gross Phase/Signal Error Detected
CHL:1 ID:0 Target ALERT: Gross Phase/Signal Error Detected
CHL:1 Drive Channel ALERT: Gross Phase/Signal Error Detected
CHL:1 RCC Channel ALERT: Unexpected Disconnect Encountered
CHL:1 ID:0 Target ALERT: Unexpected Disconnect Encountered
CHL:1 ID:0 Target ALERT: Unexpected Disconnect Encountered
CHL:1 Drive Channel ALERT: Unexpected Disconnect Encountered
CHL:1 RCC Channel ALERT: Timeout Waiting for I/O to Complete
CHL:1 ID:0 Target ALERT: Timeout Waiting for I/O to Complete
CHL:1 ID:0 Target ALERT: Timeout Waiting for I/O to Complete
CHL:1 Drive Channel ALERT: Timeout Waiting for I/O to Complete
CHL:1 RCC Channel ALERT: Parity/CRC Error Detected
CHL:1 ID:1 Host Channel ALERT: Parity/CRC Error Detected
CHL:1 Host Channel ALERT: Parity/CRC Error Detected
CHL:0 ID:0 Drive ALERT: Unexpected Drive Not Ready (00B)
CHL:0 ID:0 Drive ALERT: Drive HW Error (00B)
CHL:0 RCC Channel ALERT: Unit Attention Received
CHL:0 ID:0 Target ALERT: Unit Attention Received (10B)
CHL:0 ID:0 Target ALERT: Unit Attention Received
CHL:0 Drive Channel ALERT: Unit Attention Received
CHL:0 ID:0 Drive ALERT: Aborted Command (00B)
CHL:0 ID:0 Drive ALERT: Unexpected Sense Received (00B)
CHL:0 ID:0 Drive NOTIFY: Block Reassignment Failed - 0 (10B)
CHL:0 ID:0 Drive NOTIFY: Block Reassignment Failed - 0x0
CHL:0 RCC Channel ALERT: Data Overrun/Underrun Detected
CHL:0 ID:0 Target ALERT: Data Overrun/Underrun Detected
CHL:0 ID:0 Target ALERT: Data Overrun/Underrun Detected
CHL:0 Drive Channel ALERT: Data Overrun/Underrun Detected
CHL:0 RCC Channel ALERT: Negotiation Error Detected
CHL:0 ID:0 Target ALERT: Negotiation Error Detected
CHL:0 ID:0 Target ALERT: Negotiation Error Detected
CHL:0 Drive Channel ALERT: Negotiation Error Detected
CHL:0 RCC Channel ALERT: Invalid Status/Sense Data Received
CHL:0 ID:0 Target ALERT: Invalid Status/Sense Data Received (10B)
CHL:0 ID:0 Target ALERT: Invalid Status/Sense Data Received
CHL:0 Drive Channel ALERT: Invalid Status/Sense Data Received
CHL:0 ALERT: Redundant Loop Connection Error Detected on ID:11
CHL:0 Host Channel ALERT: Channel Failure
SMART-Slot:-1 Predictable Failure Detected-Clone Failed
SMART-Slot:-1 Predictable Failure Detected

SMART-Slot:-1 Predictable Failure Detected-Starting Clone
SMART-Slot:-1 Predictable Failure Detected (TEST)
SMART-Slot:-1 Predictable Failure Detected (TEST)
SMART-Slot:-1 Predictable Failure Detected (TEST)
SMART-Slot:-1 Predictable Failure Detected (TEST)
CHL:0 ALERT: Redundant Loop for Chl:0 Failure Detected
CHL:0 ALERT: Redundant Path for Chl:0 ID:0 Expected but Not Found
CHL:0 ID:0 ALERT: Redundant Path for Chl:0 ID:0 Failure Detected

Notification

CHL:0 ID:132 NOTICE: Media Error Recovered-0xD8001C7C
CHL:0 ID:132 NOTICE: Media Error Recovered-0xD8001C7C
CHL:0 ID:255 NOTICE: Redundant Path for Chl:0 Restored
CHL:0 NOTICE: Redundant Path for Chl:0 ID:0 Restored
CHL:0 ID:0 NOTICE: Redundant Path for Chl:0 ID:0 Restored
CHL:0 ID:0 Drive NOTICE: Drive Recovered Error Reported
CHL:0 LIP(FF B) Detected
CHL:0 Host Channel Notification: Bus Reset Issued
CHL:0 Host Channel ALERT: Bus Reset Issued
CHL:0 ID:0 Drive NOTICE: Scan Drive Successful
CHL:0 ID:0 Drive NOTIFY: Block Successfully Reassigned - 0 (10B)
CHL:0 ID:0 Drive NOTIFY: Block Successfully Reassigned - 0x0

16.1.3 General Target Events

Critical

SAF-TE Device(0)ALERT: Power Supply Failure Detected(Idx:132)
Peripheral Set 0 Device ALERT: Power Supply 0 Failure Detected
Peripheral Set 0 Device ALERT: Power Supply 0 Not Present
Peripheral Set 0 Device ALERT: Low Voltage Detected(-1073939240.5V)
SES(C0 I0) Power Supply 0: Device Not Supported!
SAF-TE Device (0)ALERT: Cooling Fan Not Installed(Idx:2)
SAF-TE Device (0)ALERT: Cooling Fan Failure Detected(Idx:2)
Peripheral Set 0 Device ALERT: Cooling fan2 Failure Detected
Peripheral Set 0 Device ALERT: Cooling fan2 Not Present
Peripheral Set 0 Device ALERT: Low Speed Detected(-1073870808 RPM)
SES(C0 I0)Cooling Fan 0:Device Not Supported!
SAF-TE Device(0)ALERT: Elevated Temperature Alert
Peripheral Device ALERT: CPU Cold Temperature Detected(0.0(C))
Peripheral Device ALERT: Elevated Temperature Alert
Peripheral Device ALERT: Elevated Temperature Alert
Peripheral Set 0 Device ALERT: Temperature Sensor 0 Failure Detected
Peripheral Set 0 Device ALERT: Temperature Sensor 0 Not Present
Peripheral Set 0 Device ALERT: Cold Detected(0(C))
SES(C0 I0)Temp Sensor 0:Device Not Supported!

Notification

SAF-TE Device(2) NOTICE: Fan Back On-Line(Idx:11)
NOTICE: FAN Module A Back On-Line(FAN0 0RPM)
NOTICE: Controller FAN 1 Back On-Line(0 RPM)
Peripheral Device NOTICE: Fan Back On-Line
Peripheral Set 1 Device NOTICE: FAN 1 Back On-Line
Peripheral Set 1 Device NOTICE: FAN 1 is Present
Peripheral Set 1 Device NOTICE: FAN 1 Back On-Line(0 RPM)
SES(C0 I0)Cooling Fan 1:Fan Back On-Line
SAF-TE Device(1) NOTICE: Temperature Back To Non-Critical LVs
Peripheral Set 1 Device NOTICE: Temperature 1 Back To Non-Critical LVs
Peripheral Set 1 Device NOTICE: Temperature 1 is Present
Peripheral Set 1 Device NOTICE: Temperature 1 Back To Non-Critical LVs(0(C))
SES(C0 I0)Temp Sensor 1:Temperature Back To Non-Critical LVs
SAF-TE Device(1) NOTICE: Power Supply Back On-Line(Idx:4)
Power Supply 0 Back-Online
Peripheral Set 2 Device NOTICE: PSU0 +12V Back On-Line
Peripheral Set 2 Device NOTICE: PSU0 +12V is Present
Peripheral Set 2 Device NOTICE: PSU0 +12V Back On-Line(5.0V)
SES(C0 I0)Power Supply 2:Power Supply Back On-Line

16.1.4 Controller Events

Critical

ALERT: +3.3V Low Voltage Detected(205.9V)
ALERT: Controller FAN 0 Low Speed Detected(0 RPM)
Controller NOTICE: Redundant Controller Firmware Updated
Controller ALERT: Redundant Controller Failure Detected
Controller SDRAM ECC Multi-bits Error Detected
Controller SDRAM ECC Single-bit Error Detected
Controller SDRAM Parity Error Detected
Controller PCI Bus Parity Error Detected
Controller ALERT: Power Supply Unstable or NVRAM Failed
BBU Absent or Failed! Correct It and Reset Ctlr to Take Effect
Controller BBU Absent or Failed!
Controller BBU Failure Detected!
Controller BBU Thermal Shutdown/Enter Sleep-Mode!

Warning

Memory Not Sufficient to Fully Support Current Config.

Notification

Board1 Cold Temperature Back To Non-Critical LVs(0(C))
+12V Upper Voltage Back within Acceptable Limits(0.0V)
+12V Lower Voltage Back within Acceptable Limits(0.0V)
Memory is Now Sufficient to Fully Support Current Config.
Controller NOTICE: NVRAM Factory Defaults Restored
Controller NOTICE: NVRAM Factory Defaults Restored
Controller NOTICE: NVRAM Factory Defaults Restored
Controller BBU is Charging!
NOTICE: Controller BBU Present!
NOTICE: Controller BBU Back On-Line!
NOTICE: Controller BBU Fully Charged!
Force Controller Write-Through on Trigger Cause

16.2

Event Description

16.2.1

Logical Drive Events

Critical

Message	LG:0 Logical Drive ALERT: CHL:0 ID:132 Drive Missing
What Happens?	The specified hard drive in the specified logical drive has is missing.
What to Do?	The drive could have been accidentally removed or drive connection problems occurred. Check drive connection and related enclosure status.

Message	LG: Logical Drive ALERT: CHL:0 ID:132 Drive Failure
What Happens?	The specified hard drive in the specified logical drive has failed
What to Do?	IF a spare is available, the subsystem will automatically start rebuild. If there is no spare, replace the faulty drive and rebuild will be automatically initiated provided that a replacement drive can be recognized by the subsystem. For example, in a SCSI-based subsystem, a manual "Scan Drive" command should be executed. In subsystems using other types of disk drives, drive swap auto-detection should be enabled on the drive channels.

Message	LG:0 Logical Drive ALERT: Rebuild Aborted
What Happens?	<p>Logical drive rebuild aborted. It could result from the one of the following reasons:</p> <ol style="list-style-type: none"> 1. The rebuild has been manually canceled by a user. 2. The replacement drive used for a rebuild may have failed during the rebuild or the drive contains immanent defects. 3. Bad blocks are encountered on another member drive during the rebuild. (RAID 1, 3, and 5 arrays) 4. System faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive rebuild again. It is best to stop host I/Os temporarily to this logical drive.

Message	LG:0 Logical Drive ALERT: Rebuild Failed
What Happens?	<p>Logical drive rebuild failed. It could result from the one of the following reasons:</p> <ol style="list-style-type: none"> 1. The rebuild has been manually canceled by a user. 2. The replacement drive used for a rebuild may have failed during the rebuild or the drive contains immanent defects. 3. Bad blocks are encountered on another member drive during the rebuild. (RAID 1, 3, and 5 arrays) 4. System faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive rebuild again. It is best to stop host I/Os temporarily to this logical drive.

Message	LG:0 Logical Drive ALERT: Parity Regeneration Aborted
What Happens?	System faults or integrity problems forced system to abandon the operation. Irrecoverable parity inconsistency may also cause a halt to the operation.
What to Do?	Check proper system working conditions. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Parity Regeneration Failed
What Happens?	System faults or integrity problems forced system to abandon the operation. Irrecoverable parity inconsistency may also cause a halt to the operation.
What to Do?	Check proper system working conditions. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Creation Aborted
What Happens?	Logical drive creation process manually aborted or that some system integrity problems forced the subsystem to abort logical drive creation process.
What to Do?	Check proper system working conditions.

Message	LG:0 Logical Drive ALERT: Creation Failed
What Happens?	Logical drive creation process failed when one or more drive members failed or have been accidentally removed. Drive abnormality may also be the cause of the problems.
What to Do?	Check proper system working conditions and drive connection. Replace the faulty drives and repeat the creation process.

Message	LG:0 Logical Drive ALERT: Initialization Failed
What Happens?	<p>Logical drive initialization failed. It could result from one of the following reasons:</p> <ol style="list-style-type: none"> 1. One of the disk drives used for a logical drive might have failed during rebuild. 2. Media errors are encountered on another member drive during the initialization process. 3. Drive hardware faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive initialization again.

Message	LG:0 Logical Drive ALERT: Expansion Aborted
What Happens?	<p>Logical drive expansion failed. It could result from one of the following reasons:</p> <ol style="list-style-type: none"> 1. The expansion has been canceled by a user. 2. The drive used for expansion might have failed during expansion. 3. Media errors are encountered on another member drive during the expansion process. 4. One member drive has failed. 5. Drive hardware faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive expansion again.

Message	LG:0 Logical Drive ALERT: Expansion Failed
What Happens?	<p>Logical drive initialization failed. It could result from one of the following reasons:</p> <ol style="list-style-type: none"> 1. The drive(s) used for expansion might have failed during expansion. 2. Media errors are encountered on another member drive during the expansion process. 3. One member drive has failed. 4. Drive hardware faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive expansion again.

Message	LG:0 ALERT: CHL:0 ID:132 Media Scan Failed
What Happens?	Media Scan failed to carry on with scanning drive(s) for certain reasons, e.g., a logical drive or a disk drive has fatally or physically failed.
What to Do?	Carefully identify and replace the faulty drive if rebuild is possible.

Message	LG:0 ALERT: CHL:0 ID:132 Media Scan Aborted
What Happens?	Media Scan operation is aborted for certain reasons, e.g., it has been manually canceled.
What to Do?	Check disk drive and system operation status.

Message	LG:0 Logical Drive ALERT: CHL:0 ID:132 Clone Failed
What Happens?	<p>Drive cloning failed. It could result from one of the following reasons:</p> <ol style="list-style-type: none"> 1. Cloning has been manually canceled by a user. 2. The drive used for cloning might have failed during the cloning process. 3. Bad blocks are encountered on another member drive during the cloning process. 4. The source drive has failed. 5. System faults occurred.
What to Do?	Carefully identify and replace the faulty drive.

Message	LG:0 Logical Drive ALERT: CHL:0 ID:132 Clone Aborted
What Happens?	<p>Logical drive cloning aborted. It could result from one of the following reasons:</p> <ol style="list-style-type: none"> 1. The cloning process has been manually canceled by a user. 2. The drive used for cloning might have failed during cloning. 3. Media errors are encountered on another member drive during the cloning process. 4. The source drive has failed. 5. System has been reset or drive hardware faults occurred.
What to Do?	Carefully identify and replace the faulty drive.

Message	LG:0 Logical Drive ALERT: Logical Drive Bad Block Table FULL
What Happens?	Available bad block table entries full. Media Scan or host access has discovered and marked too many damaged sectors.
What to Do?	The target disk drive may have become physically unreliable. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Logical Drive Bad Block Table BAD
What Happens?	System failed to generate a bad block table. Logical drive may have generated serious integrity problems.
What to Do?	The members of a logical drive may have become physically unreliable. Locate, replace, and replace the logical drive if possible. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Logical Drive On-Line Init Table Bad
What Happens?	Serious system faults might have occurred, unstable voltage or damaged components, and forced the initialization to halt.
What to Do?	Check proper system working conditions. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Bad Data Block Marked 000000084
What Happens?	<p>Media errors caused data inconsistency; e.g., concurrent occurrences of media errors on more than one RAID5 member drive. The affected data blocks # are marked bad. The situation may result from the following:</p> <ol style="list-style-type: none"> 1. One member drive failed, and media errors are found on another member drive. 2. Media errors occurred consequentially on two member drives.
What to Do?	The data affected by data blocks failure will be lost. The host computer will return media error messages if host accesses fall within the affected blocks.

Message	LG:0 Logical Drive ALERT: UNPROTECTED Block Marked 000000084
What Happens?	Media errors are found in an unprotected array, e.g., RAID0. The affected data blocks will be marked bad.
What to Do?	The data affected by data blocks failure will be lost. The host computer will return media error messages if host access falls within the affected blocks. Restoring data from a backup source is recommended.

Message	LG:0 Logical Drive NOTIFY: Bad Data Block encountered 000000084
What Happens?	Data blocks previously marked bad are re-encountered during host I/O access or certain operations such as Media Scan.
What to Do?	The data affected by data blocks failure is lost. Restoring data from a backup source is recommended. The host computer will return media error messages if host access falls within the affected blocks.

Message	LG:0 ALERT: Inconsistent Parity Encountered Block 000000084
What Happens?	Parity inconsistency encountered on block address ____
What to Do?	It is recommended to perform the "Regenerate Parity" function to correct the parity faults.

Notification:

Message	LG:0 Logical Drive NOTICE: Continue Rebuild Operation
What Happens?	The target logical drive has been stored to its previous one-drive-failed status, and the rebuild operation is automatically resumed by firmware. This may occur when the system is powered off or reset during the rebuild process or that media errors have been discovered and corrected during the rebuild.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Rebuild
What Happens?	The rebuild process has begun.
What to Do?	This is the message displayed when a stand-by spare is available when a member drive fails or when a faulty drive is physically replaced. The subsystem should automatically detect a drive for rebuild if the hot-swap detection mechanism has been enabled.

Message	LG:0 Logical Drive NOTICE: Rebuild of Logical Drive Completed
What Happens?	The subsystem has successfully rebuilt a logical drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Continue Parity Regeneration Operation
What Happens?	Start
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Parity Regeneration
What Happens?	Start regenerating parity of a logical drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Parity Regeneration of Logical Drive Completed
What Happens?	The parity regeneration process on logical drive _ is completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Creation
What Happens?	Logical drive creation process started. Note that the creation and initialization processes are separated. Creation of a logical drive only takes a second while the initialization may take hours with the arrays containing large capacity drives.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting On-Line Initialization
What Happens?	Creation process is completed, initialization process has begun to generate parity sectors and readying member drives. The On-line initialization mode takes a longer time to complete and allows the logical drive to receive host I/Os immediately if appropriate host ID/LUN mapping has been applied.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: On-Line Initialization of Logical Drive Completed
What Happens?	Logical drive on-line initialization completed
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Off-Line Initialization
What Happens?	Creation process is completed, initialization process has begun to generate parity sectors and readying member drives. The Off-line initialization mode takes place immediately after the creation of a logical drive, and does not allow host I/Os to be distributed to the logical drive at the same time.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Off-Line Initialization of Logical Drive Completed
What Happens?	Logical drive off-line initialization completed. The logical drive is now ready for host ID/LUN mapping, and then receiving host I/Os.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting On-Line Expansion
What Happens?	Start expanding the logical drive's added or free capacity (previously not included) when system find appropriate time after selecting to expand its capacity. The On-Line mode allows the added capacity of a logical drive to appear immediately as a usable partition and ready for host I/Os. This partition may require being associated with a separate ID/LUN on the host bus. Response to host I/Os will be slower because the expansion process requires distributing data blocks and parity data across new drive sectors.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: On-Line Expansion of Logical Drive Completed
What Happens?	Logical drive on-line expansion completed The added capacity will appear as a new partition, and this partition may require being associated with a separate ID/LUN on the host bus. Adjustments should also be made through the host OS volume management utility.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Off-Line Expansion
What Happens?	Start expanding the logical drive's added or free capacity (previously not included) when system find appropriate time after selecting to expand its capacity. The Off-Line mode does not allow the added capacity of a logical drive to be immediately available for host I/Os. This partition can only be associated with a separate ID/LUN on the host bus after the expansion process is completed. Access to the data already stored on the logical drive is not affected.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Continue Add Drive Operation
What Happens?	The target logical drive has been restored to its previous status, e.g., power-off or media errors, and the add drive operation can continue.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Add Drive Operation
What Happens?	Add drive expansion process started
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Add Drive Operation Paused
What Happens?	The add drive expansion process is halted by: <ol style="list-style-type: none"> 1. Logical drive expansion is canceled by a user. 2. One of the member drives failed during logical drive expansion. 3. Media errors are found on one or more of the logical drive members. 4. System faults have occurred.
What to Do?	Carefully identify and replace the faulty drive and correct system faults. The add drive expansion should continue once the faults are corrected.

Message	LG:0 Logical Drive NOTICE: Add Drive to Logical Drive Completed
What Happens?	The add drive expansion process is completed.
What to Do?	Press ESC to clear the message. Please remember to rearrange the added capacity which will appear as a new partition volume.

Message	NOTICE: CHL:0 ID:132 Starting Media Scan
What Happens?	Media Scan is manually or automatically started by the preset Task Schedules on a specific disk drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 NOTICE: CHL:0 ID:132 Starting Media Scan
What Happens?	Media Scan is manually or automatically started by the preset Task Schedules on a specific member of Logical Drive 0.
What to Do?	Press ESC to clear the message.

Message	NOTICE: Media Scan of CHL:0 ID:132 Completed
What Happens?	Media Scan on a specific disk drive is completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 NOTICE: Media Scan of CHL:0 ID:132 Completed
What Happens?	Media Scan on a specific member of Logical Drive 0 is completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Continue Clone Operation
What Happens?	The target logical drive has been restored to its previous status, e.g., a failed drive is restored, and the cloning process can continue.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Starting Clone
What Happens?	Cloning process started on one of the members of Logical Drive 0
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Copy and Replace Completed
What Happens?	Copy and replace is completed with a specific member of Logical Drive 0. The original member drive is replaced by another disk drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Clone Completed
What Happens?	Cloning is completed with a specific member of Logical Drive 0, whether it is replaced (Copy and Replace) or cloned (Perpetual Clone) to a spare drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTIFY: Bad Data Block Recovered 000000084
What Happens?	Host writes fall onto the data blocks previously marked bad and overwrite the drive sectors.
What to Do?	Press ESC to clear the message.

16.2.2 Channel and Individual Drive Events

Critical:

Message	CHL:0 ID:132 Target ALERT: Media Error Unrecoverable-0xD8001C7C
What Happens?	Drive surface media errors found and after rewrite the drive sectors are determined as physically damaged and unusable.
What to Do?	Data will be reassigned to other drive sectors. Host writes should continue without problems.

Message	CHL:0 Fatal Error (0)
What Happens?	Channel 0 has fatally failed.
What to Do?	Contact your system vendor for help.

Warning

Message	CHL:1 ID:0 Target ALERT: Unexpected Select Timeout
What Happens?	Drive target select timeout. The specified hard drive cannot be selected by the RAID controller/subsystem. Whether the drive has been removed, or the cabling/termination/canister is out of order.
What to Do?	This could occur if a disk drive is not properly installed or accidentally removed, or when cable links become unreliable. Check cabling and drive installation, and contact your RAID system supplier.

Message	CHL:1 Drive Channel ALERT: Unexpected Select Timeout
What Happens?	Drive Channel target select timeout. The specified drive channel cannot be selected by the RAID controller/subsystem. Faults might have occurred with the cabling/termination/internal data paths.
What to Do?	Check cabling and drive installation, and contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Gross Phase/Signal Error Detected
What Happens?	Phase/signal abnormality detected on the specific RCC channel.
What to Do?	Rare occurrence of phase/signal error could be recovered by firmware retry. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Gross Phase/Signal Error Detected
What Happens?	Phase/signal abnormality detected with a specific target disk drive.
What to Do?	Check proper installation of disk drives and contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Gross Phase/Signal Error Detected
What Happens?	Phase/signal abnormality detected on the specific drive channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Unexpected Disconnect Encountered
What Happens?	Unexpected disconnect detected on the specific RCC channel.
What to Do?	Check cabling/termination and canister connections, and contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Unexpected Disconnect Encountered
What Happens?	Unexpected disconnect detected on the specific target disk drive.
What to Do?	Check cabling/termination and canister installation, and contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Unexpected Disconnect Encountered
What Happens?	Unexpected disconnect detected on the specific drive channel.
What to Do?	Check cabling/termination and canister connections, and contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Timeout Waiting for I/O to Complete
What Happens?	I/O timeout on specific RCC channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Timeout Waiting for I/O to Complete
What Happens?	Drive-side target I/O timeout. Possible drive-side cabling/termination and canister connection abnormal or drive malfunctioning.
What to Do?	Check drive-side cabling/termination/canister/disk drive connections and contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Timeout Waiting for I/O to Complete
What Happens?	I/O timeout on specific drive channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Parity/CRC Error Detected
What Happens?	RCC channel parity or CRC errors detected.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Host Channel ALERT: Parity/CRC Error Detected
What Happens?	Parity or CRC errors detected through host ID, CHL_ ID_.
What to Do?	Occurrences of these faults usually can be negotiated between RAID subsystems and application servers. If occurrences become frequent, contact your RAID system supplier.

Message	CHL:_ Host Channel ALERT: Parity/CRC Error Detected
What Happens?	Host channel parity or CRC errors detected.
What to Do?	Parity and CRC faults can often be negotiated between application servers and RAID subsystems. Contact your RAID system supplier if this occurs too frequently.

Message	CHL:_ ID:_ Drive ALERT: Unexpected Drive Not Ready (00B)
What Happens?	Unexpected Drive Not Ready returned from a disk drive on CHL:_ ID:_.
What to Do?	Check proper installation of disk drives or cabling connection. This event often occurs when hot-swapping disk drives and may be caused by signal glitches. If the disk drive fails to respond, the subsystem should consider the disk drive as a failed drive when the timeout threshold is reached. Contact your RAID system supplier.

Message	CHL:_ ID:_ Drive ALERT: Drive HW Error (00B)
What Happens?	Drive hardware error returned from a disk drive on CHL:_ ID:_.
What to Do?	The target disk drive may have failed. Contact your RAID system supplier for a replacement.

Message	CHL:_ RCC Channel ALERT: Unit Attention Received
What Happens?	Unit attention received from RCC channel CHL:_.
What to Do?	Rare occurrences of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Unit Attention Received (10B)
What Happens?	Drive-side target unit attention receive on a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Unit Attention Received (10B)
What Happens?	Drive-side target unit attention receive on a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Check disk drive connection and contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Unit Attention Received (10B)
What Happens?	Drive-side target unit attention receive on a drive channel CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Check proper installation and contact your RAID system supplier.

Message	CHL:_ ID:_ Drive ALERT: Aborted Command (00B)
What Happens?	Aborted command reported from a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Drive ALERT: Unexpected Sense Received (00B)
What Happens?	Drive-side target unexpected sense received on a disk drive CHL:_ ID:_.
What to Do?	Check drive-side cabling/installation/working conditions. This may result from bad signal quality or poor connection. Contact your RAID system supplier.

Message	CHL:_ ID:_ Drive NOTIFY: Block Reassignment Failed – 0(10B)
What Happens?	Data regeneration and reassignment failed. The Bad Block Table is full. Too many bad blocks have been found on the disk drive. The adjacent data blocks on other member drives might have also been damaged.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Data Overrun/Underrun Detected
What Happens?	Data overrun/underrun errors detected on the RCC channel CHL:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Data Overrun/Underrun Detected
What Happens?	Data overrun/underrun errors detected on a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Data Overrun/Underrun Detected
What Happens?	Data overrun/underrun errors detected on the drive channel CHL:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Negotiation Error Detected
What Happens?	Negotiation errors occurred on the RCC channel CHL:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Negotiation Error Detected
What Happens?	Negotiation errors occurred with the communications with a disk drive CHL:_ ID:_ . The event could also occur with drive target data bus sync/wide negotiation abnormality.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Negotiation Error Detected
What Happens?	Negotiation errors occurred with the communications with a drive channel CHL:_ . The event could also occur with drive target data bus sync/wide negotiation abnormality.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Invalid Status/Sense Data Received
What Happens?	Invalid status/sense data received on the RCC channel CHL:_
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target ALERT: Invalid Status/Sense Data Received (10B)
What Happens?	Invalid status/sense data received by a disk drive CHL:_ ID:_
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Invalid Status/Sense Data Received
What Happens?	Invalid status/sense data received on the drive channel CHL:_
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ALERT: Redundant Loop Connection Error Detected on ID:_
What Happens?	One of the dual loop members might have failed or been disconnected. Make sure all cabling are properly connected and topological configurations properly set.
What to Do?	Check cabling and channel bus working conditions. If the problem does not result from cabling problems, contact your RAID system supplier.

Message	CHL:_ Host Channel ALERT: Channel Failure
What Happens?	Host channel CHL:_ link failure or disconnection occurred.
What to Do?	Check proper cabling connection and host HBA, link connection devices, etc. Contact your RAID system supplier.

Message	SMART-Slot:_ Predictable Failure Detected-Clone Failed
What Happens?	Hard Drive(s) report SMART-detected defects. A spare drive is commenced to clone or replace the member drive showing defects. The cloning operation failed. The member drive suspected of errors might have failed, or the clone process has been interrupted, e.g., yet another member has failed or the spare drive used for cloning has shown immanent faults.
What to Do?	Carefully check drive status and replace the unstable/failed drive. Contact your RAID system supplier.

Message	SMART-Slot:_ Predictable Failure Detected
What Happens?	Hard drive slot:_ reported SMART-detected errors, e.g., abnormal rotation speed, seek time delay, etc.
What to Do?	Carefully check drive status and replace the unstable/failed drive. Contact your RAID system supplier.

Message	SMART-Slot:_ Predictable Failure Detected-Starting Clone
What Happens?	One or more hard Drive(s) reports SMART-detected defects. Starts using a pre-defined spare drive to clone the member drive suspected of inherent errors. This happens when SMART detection is enabled with the “Copy & Replace” or “Perpetual Clone” reaction schemes.
What to Do?	Press ESC to clear the message. After the cloning process, you may consider replacing the defective drive with a good one and configure it as a spare drive in case drive failure might happen in the future.

Message	SMART-Slot:_ Predictable Failure Detected (TEST)
What Happens?	This results from a manually executed test command. This ensures your disk drive supports SMART detection and report functionality.
What to Do?	Press ESC to clear the message.

Message	CHL:_ ALERT: Redundant Loop for CHL:_ Failure Detected
What Happens?	One of the dual loop members may have failed or been disconnected. Make sure all channels are properly connected and topological configurations properly set.
What to Do?	Check cabling and system installation. Contact your supplier if no connection errors could be found.

Message	CHL:_ ALERT: Redundant Path for CHL:_ ID:_ Expected but Not Found
What Happens?	One of the dual loop members may have failed or been disconnected. This event is regenerated in the event when the one of dual loop connections to an expansion enclosure can not be found after power-up. Make sure all channels are properly connected and topological configurations properly set.
What to Do?	Check cabling and system installation. A manual “Issue LIP” command may restore the connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ ID:_ ALERT: Redundant Loop for CHL:_ ID:_ Failure Detected
What Happens?	One of the dual loop members connecting device CHL:_ ID:_ may have failed or been disconnected. Make sure all channels are properly connected and topological configurations properly set.
What to Do?	Check cabling and system installation. Contact your supplier if no connection errors could be found.

Notification

Message	CHL:_ ID:_ NOTICE: Media Error Recovered-0xD8001C7C
What Happens?	Data once stored on bad drive sectors has been successfully re-conducted from adjacent data blocks/stripes onto the original drive blocks or distributed onto other healthy sectors. ** Bosch firmware translates SATA error codes into SCSI standard sense keys and sense codes. That's the error codes at the end of message.
What to Do?	Press ESC to clear the message.

Message	CHL:_ ID:_ NOTICE: Redundant Path for CHL:_ Restored
What Happens?	Redundant path, one of the drive-side loop members, for CHL:_ ID:_ is restored.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ NOTICE: Redundant Path for CHL:_ ID:_ Restored
What Happens?	The redundant path (CHL:_) connecting drive (CHL:_ ID:_) is restored.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ ID:_ Drive NOTICE: Drive Recovered Error Reported
What Happens?	Hard drive unrecoverable media error reported. A bad block is encountered in the specified hard drive. The RAID controller will ask the hard drive to retry.
What to Do?	Press [ESC] to clear the message.

Message	CHL:_ LIP(FF B) Detected
What Happens?	A LIP is issued on CHL:_.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ Host Channel Notification: Bus Reset Issued.
What Happens?	Bus reset on the CHL:_ host channel.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ ID:_ Drive NOTICE: Scan Drive Successful
What Happens?	A new drive or a replacement drive has been successfully scanned/recognized by the subsystem.
What to Do?	Press ESC to clear the message.

Message	CHL:_ ID:_ Drive NOTIFY: Block Successfully Reassigned – 0 (10B)
What Happens?	Writes conducted, retried, and media errors were discovered. Writes will then be re-directed to other drive sectors.
What to Do?	Press ESC to clear the message. If this message becomes frequent, the target disk drive might have developed serious media errors.

Message	CHL:_ ID_ Drive #SEV#: Media Error Encountered
What Happens?	Media errors are encountered in a hard drive.
What to Do?	Press ESC to clear the message. If this message becomes frequent, the target disk drive might have developed serious media errors.

16.2.3 General Target Events

Critical

Message	SAF-TE Device(0) ALERT: Power Supply Failure Detected(Idx:132)
What Happens?	Power supply failure detected through SAF-TE monitoring device.
What to Do?	Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Power Supply 0 Failure Detected
What Happens?	Power supply failure detected through an I2C serial bus.
What to Do?	Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Power Supply 0 Not Present
What Happens?	No power supply module is installed in the expected module slot.
What to Do?	Check proper module installation. If an installed module can not be detected, contact your RAID supplier.

Message	Peripheral Set 0 Device ALERT: Low Voltage Detected (-1073939240.5V)
What Happens?	Low voltage detected from the power supply. Power supply may have become unstable.
What to Do?	Check proper module installation and contact your RAID supplier for a replacement module.

Message	SES(C0 I0)Power Supply 0:Device Not Supported!
What Happens?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Message	SAF-TE Device(0) ALERT: Cooling Fan Not Installed(Idx:2)
What Happens?	No cooling fan module is installed in the expected module slot.
What to Do?	Check proper module installation. If an installed module can not be detected, contact your RAID supplier.

Message	SAF-TE Device(0) ALERT: Cooling Fan Failure Detected(Idx:2)
What Happens?	Cooling fan failure detected through SAF-TE monitoring device.
What to Do?	Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Cooling Fan 2 Failure Detected
What Happens?	Cooling fan failure detected through an I2C serial bus.
What to Do?	Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Cooling Fan 2 Note Present
What Happens?	No cooling fan module is installed in the expected module slot.
What to Do?	Check proper module installation or contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Low Speed Detected(-1073870808 RPM)
What Happens?	Low rotation speed detected on cooling fan. This is a fault detected through an I2C serial bus.
What to Do?	Check proper module installation or contact your RAID supplier for a replacement module.

Message	SES(C0 I0)Cooling Fan 0:Device Not Supported!
What Happens?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Message	SAF-TE Device(0) ALERT: Elevated Temperature Alert
What Happens?	Elevated temperature detected through SAF-TE monitoring device. Improper installation site condition or cooling fan failure might have caused the system temperature to rise.
What to Do?	Check proper cooling and cooling fan working condition. Contact your RAID system supplier.

Message	Peripheral Device ALERT: CPU Cold Temperature Detected(._C))
What Happens?	Temperature below the lower CPU threshold detected. Improper installation site condition might have caused the situation.
What to Do?	Correct your installation site condition. Contact your RAID system supplier.

Message	Peripheral Device ALERT: Elevated Temperature Alert
What Happens?	High temperature threshold violated and detected through an I2C serial bus. Improper installation site condition or cooling fan failure might have caused the situation.
What to Do?	Correct your installation site condition and cooling fan failure. Contact your RAID system supplier if a replacement is necessary.

Message	Peripheral Set 0 Device ALERT: Temperature Sensor 0 Failure Detected
What Happens?	Temperature sensor failure reported through an I2C serial bus.
What to Do?	Contact your RAID system supplier.

Message	Peripheral Set 0 Device ALERT: Temperature Sensor 0 Not Present
What Happens?	Temperature sensor failed or serial bus configuration/connection faults occurred.
What to Do?	Contact your RAID system supplier.

Message	Peripheral Set 0 Device ALERT: Cold Detected
What Happens?	Temperature dropped below the lower system threshold detected. Improper installation site condition might have caused the situation.
What to Do?	Correct your installation site condition. Contact your RAID system supplier.

Message	SES(C0 I0)Temp Sensor 0: Device Note Supported!
What Happens?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Notification:

Message	SAF-TE Device(2) NOTICE: Fan Back On-Line(Idx:11)
What Happens?	A once missing or failed cooling fan is restored. This message is reported through the SAF_TE monitoring device.
What to Do?	Press ESC to clear this message.

Message	NOTICE: Fan Module A Back On-Line(FAN0_RPM)
What Happens?	A once missing or failed cooling fan is restored. This message is reported through an I2C serial bus with RPM reading.
What to Do?	Press ESC to clear this message.

Message	Peripheral Device NOTICE: Fan Back On-Line
What Happens?	An I2C serial bus message indicating a once missing or failed cooling fan is restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: FAN 1 Back On-Line
What Happens?	An individual cooling fan once missing or failed is restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: FAN 1 is Present
What Happens?	An individual cooling fan once missing or failed is restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: FAN 1 Back On-Line(_RPM)
What Happens?	An individual cooling fan once missing or failed is restored, and the PRM reading displays. This applies to enclosure implementations using cooling fans capable of reporting rotation speed.
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0) Cooling Fan 1:Fan Back On-Line
What Happens?	A cooling fan once missing or failed is restored. This message is reported through an SES monitoring device.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: FAN 1 Back On-Line
What Happens?	An individual cooling fan once missing or failed is restored.
What to Do?	Press ESC to clear this message.

Message	SAF-TE Device(1) NOTICE: Temperature Back To Non-Critical LVs
What Happens?	Operating temperature back within normal temperature range.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: Temperature 1 Back To Non-Critical LVs
What Happens?	Temperature sensor 1 reports operating temperature is now within normal temperature range.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: Temperature 1 is Present
What Happens?	Temperature sensor 1 once missing and is now present on the serial bus.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: Temperature 1 Back To Non-Critical LVs (_C)
What Happens?	Temperature sensor 1 reports operating temperature is now within normal temperature range and the Celsius reading displays.
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0) Temp Sensor 1: Temperature Back to Non Critical LVs
What Happens?	Temperature sensor 1 reports operating temperature is now within normal temperature range.
What to Do?	Press ESC to clear this message.

Message	SAF-TE Device(1) NOTICE: Power Supply Back On-Line (Idx:4)
What Happens?	A power supply once missing or failed is now restored. This event is reported through the SAF-TE monitoring interface.
What to Do?	Press ESC to clear this message.

Message	Power Supply 0 Back-Online
What Happens?	Power supply module 0 is back online
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: PSU0 +12V Back On-Line
What Happens?	Power supply module 0 restored normal +12V voltage range
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: PSU0 +12V is Present
What Happens?	Power supply module 0 restored normal +12V voltage range
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0)Power Supply 2: Power Supply Back On-Line
What Happens?	Power supply module 2 once missing or failed is now restored.
What to Do?	Press ESC to clear this message.

16.2.4

Controller Events

Critical

Message	ALERT: +3.3V Low Voltage Detected(_V)
What Happens?	The detected +3.3V voltage source is lower than the preset voltage threshold.
What to Do?	Check power supply working condition, voltage threshold settings, and contact your RAID system supplier if necessary.

Message	ALERT: Controller FAN 0 Low Speed Detected (_RPM)
What Happens?	This only applies to controller fans in the SentinelRAID series. Low rotation speed detected.
What to Do?	Contact your RAID system supplier for a replacement.

Message	Controller NOTICE: Redundant Controller Firmware Updated
What Happens?	Firmware updated on a redundant controller RAID subsystem. Resetting the subsystem may be required for the new firmware to take effect.
What to Do?	Updating firmware on a complex system configuration is recommended for experienced engineers only. Unless for bug fixing, updating firmware is not always necessary.

Message	Controller ALERT: Redundant Controller Failure Detected
What Happens?	One RAID controller in a controller pair failed. This message indicates that its partner has taken the workload and managed the failover operation.
What to Do?	Contact your RAID system supplier for a replacement controller and restore the dual-active configuration.

Message	Controller SDRAM ECC Multi-bits Errors Detected
What Happens?	Memory multi-bits errors.
What to Do?	Contact your RAID system supplier to return the subsystem/RAID controller for repair.

Message	Controller SDRAM ECC Single-bit Errors Detected
What Happens?	Memory single-bit errors.
What to Do?	Subsystem firmware should be able to manage the single-bit errors. Contact your RAID system supplier if this message becomes frequent.

Message	Controller SDRAM Parity Errors Detected
What Happens?	Memory parity errors.
What to Do?	Subsystem firmware should be able to manage the parity errors in memory. Contact your RAID system supplier if this message becomes frequent.

Message	Controller PCI Bus Parity Errors Detected
What Happens?	Serious system faults have occurred.
What to Do?	Stop host I/Os immediately to prevent further data dis-integrity and contact your RAID system supplier.

Message	Controller ALERT: Power Supply Unstable or NVRAM Failed
What Happens?	Power supply voltage fluctuating or internal system faults have occurred. Different system faults may trigger this event, e.g., memory pins contact problems or incongruent firmware versions, etc.
What to Do?	Check proper system operation conditions and contact your RAID system supplier.

Message	BBU Absent or Failed! Correct It and Reset Ctrlr to Take Effect
What Happens?	BBU (battery cell pack) has failed or is accidentally removed. When a new module is installed, reset the subsystem for the configuration to take effect.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a replacement module.

Message	Controller BBU Absent or Failed!
What Happens?	BBU (battery cell pack) has failed or is accidentally removed. When a new module is installed, reset the subsystem for the configuration to take effect.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a replacement module.

Message	Controller BBU Failure Detected!
What Happens?	BBU (battery cell pack) has failed or is accidentally removed. When a new module is installed, reset the subsystem for the configuration to take effect.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a replacement module.

Message	Controller BBU Thermal Shutdown/Enter Sleep-Mode!
What Happens?	<p>The following conditions will cause the charger circuits to enter a low-power and self-protection state:</p> <ol style="list-style-type: none"> 1. The temperature sensor on the charger circuit reports elevated temperature reading (≥ 45 degree Celsius). 2. The BBU (battery cell pack) has been charged for over 7 hours. The BBU charger will enter a timer fault state.
What to Do?	<ol style="list-style-type: none"> 1. Check proper ventilation within the subsystem. You may also check the readings from other sensors within the enclosure. Air-flow might have been disrupted by the absence of one or several major modules or the failure of a cooling fan. Once the thermal condition is improved, charging will resume automatically. 2. If a new battery module has been charged for over seven (7) hours and this event is issued, you may remove and re-install the battery module. An empty battery module may take more than 7 hours to be fully charge. There is a timer embedded with the charger, doing so can reset the timer. Charging will resume automatically.

Warning

Message	Memory Not Sufficient to Fully Support Current Config.
What Happens?	The installed memory size does not support current configuration. Try using a DIMM module of a larger size.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a certified module.

Notification

Message	Board1 Cold Temperature Back to Non-Critical LVs(_C))
What Happens?	Board1 (usually the RAID controller main board) ventilation condition is restored to normal.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	+12V Upper Voltage Back within Acceptable Limits (__V)
What Happens?	+12V voltage dropped back within preset thresholds.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	+12V Lower Voltage Back within Acceptable Limits (__V)
What Happens?	+12V voltage dropped back within preset thresholds.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	Memory is Now Sufficient to Fully Support Current Config.
What Happens?	Memory of a sufficient capacity is installed and the subsystem should work properly.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	Controller NOTICE: NVRAM Factory Defaults Restored
What Happens?	The factory defaults for configuration profile is restored.
What to Do?	Press ESC to clear the message.

Message	Controller BBU is Charging!
What Happens?	Once a new BBU is installed or BBU charge drops below a preconfigured level, the subsystem will start charging its batteries.
What to Do?	Press ESC to clear the message. All batteries have a limited life expectancy. If this message appears too frequently, the BBU may have gradually lost its ability to hold charge.

Message	+12V Upper Voltage Back within Acceptable Limits (_ _ V)
What Happens?	+12V voltage dropped back within preset thresholds.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	NOTICE: Controller BBU Present!
What Happens?	A BBU once missing or removed is now electrically connected to the subsystem.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	NOTICE: Controller BBU Back On-Line!
What Happens?	A BBU once missing or removed is now restored to its normal working condition.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	NOTICE: Controller BBU Fully Charged!
What Happens?	A BBU is fully charged.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	Force Controller Write-Through on Trigger Cause
What Happens?	Preset conditions, e.g., elevated temperature or component failure, has forced the subsystem to adopt a more conservative operating mode by disabling the Write-Back caching.
What to Do?	Check proper working condition and correct system faults. Press ESC to clear the message.

Index

Symbols

(Multi-Level RAID) 19
 , Device Support for Removable Media 134
 /O sorting 191

Numerics

2nd drive failure 187

A

Abort clone 178
 Adaptive Write Policy 191
 Add Drive expansion 163
 Add Drive Procedure 163
 Added Capacity 161
 added capacity 163
 adding a schedule 154
 Adding a SCSI Channel's ID 70
 Adding a SCSI Channel's ID LCD 70
 Adding a SCSI ID terminal 119
 Adjust LCD Contrast 77, 123
 Advanced Configurations 131, 143, 176
 Advanced Features 189
 Advanced Maintenance 198
 Alert 200
 All undefined LUNs 134
 ANSI Mode 82
 ANSI+Color Mode 82
 ANSI/VT-100 198
 Array Expansion 161
 Assign Spare Drives terminal 96
 Assigning a Logical Drive Name 65
 Assigning a Logical Drive Name LCD 65
 Assigning Logical Drive Name terminal 109
 Assigning Spare Drive 61
 Assigning Spare Drive LCD 61
 Assigning Spare Drive terminal 105
 Auto cache flush 189
 Auto recovery from logical drive failure 187
 Auto Shutdown 145
 Auto-Assign 152
 Auto-Assign Global Spare Drive 152
 Auto-Detect Failure Drive Swap Check Time 109
 Auto-Flush 145
 Automatic engagement of replacement controller 193
 Automatic Rebuild 17
 Automatic Shutdown 194
 Auto-rebuild 187
 Auto-rebuild onto spare 187
 Auto-scan of replacement drive 187
 Auto-Shutdown 145
 AV Applications 184

B

Background Download 197
 Background firmware download 187
 Bad block auto-reassignment 194
 Bad Block Handling 189
 Basic RAID Management 186
 battery 193
 battery charge 83
 BBU low 143
 beeper mute 74, 121

block LBA 67
 block number 51, 94
 Boot Record 197, 198

C

cache parameters 48, 91
 Cache Status 82
 Caching Mode Dynamic Switch 189
 Caching Optimization 191
 Caching Parameters 48, 49, 50, 91, 143
 Caching Parameters LCD 48
 Change Password 74
 Change Password LCD 74
 Change Password terminal 121
 Changing Password 74
 changing password 74, 121
 Changing the Password terminal 122
 Channel failure protection 20
 Choosing Member Drives 52
 Choosing Member Drives LCD 52
 Choosing Member Drives terminal 95
 Choosing RAID Level terminal 95
 clone drive 177, 179
 Clone Failing Drive 114, 176, 194
 clone failing drive 176
 Clone plus Replace 180
 cloning progress 183
 Co-existing spares 187
 COM port 197
 Communication Parameters 39, 40
 Concurrent Rebuild 18
 Concurrent rebuild 187
 Configuration on Disk 189
 Configuration on Drives 187
 Continuous Sync 50, 93, 191
 Controller Date and Time 79, 128
 Controller Initialization 155
 Controller Maintenance 75, 198
 Controller Name 78, 82
 Controller Name LCD 78
 Controller Name terminal 125
 controller naming 125
 Controller Parameters 30, 78
 Controller Parameters LCD 78
 Controller Parameters terminal 125
 Controller Peripheral Device Configuration 147
 controller reset 75, 123
 Controller Self-Monitoring 147
 controller temperature sensors 196
 Controller Unique Identifier 78
 Controller Unique Identifier LCD 78
 Controller Unique Identifier terminal 127
 Controller Voltage and Temperature 46
 LCD 46
 Controller Voltage and Temperature terminal 88
 controller voltage monitors 195
 Controller/Logical Drive Shutdown 190
 Copy and Replace Drives expansion 166
 Copy and Replace Procedure 166
 Creating a Logical Drive 52

- LCD 52
- Creating a Logical Drive terminal 95
- Creating a Logical Volume LCD 56
- Creating a Logical Volume terminal 99
- creating a task schedule 154
- Cursor Bar 82
- CurSynClk 88
- CurWid 88
- Custom Inquiry Serial Number 196
- Cylinder 137
- Cylinder Head Sector Mapping 137
- Cylinder Ranges 138
- D**
- Data Integrity 151
- Data Rate 120
- Date & Time 82
- Date and Time 79, 128
- Dedicated Spare Drive 186
- default setting restoring 72
- DefSynClk 87
- DefWid 87
- deleting
 - partition of logical drive 64
 - SCSI channel ID 71
 - spare drive, global or local 114
- Deleting a Logical Drive 63
- Deleting a Logical Drive LCD 63
- Deleting a Partition of a Logical Drive 64
- Deleting a Partition of a Logical Drive LCD 64
- Deleting a Partition of a Logical Drive terminal 108
- deleting a schedule 154
- Deleting a SCSI Channel's ID 71
- Deleting a SCSI Channel's ID LCD 71
- Deleting a SCSI ID terminal 119
- deleting logical drive 108
- Deleting Logical Drive terminal 108
- deleting partitions of logical drive 108
- deleting SCSI ID 119
- Deleting Spare Drive 69
- Deleting Spare Drive LCD 69
- Deleting Spare Drive terminal 114
- deleting spare drive, global or local 69
- Detect Clone plus Replace 183
- Detect Only 182
- Detect Only SMART 179
- Detect Perpetual Clone 183
- Device Check Time 149
- Device Qualifier Support 134
- Device Support for Removable Media 135
- Device Type 134
- Disabling Password 75
- Disabling Password LCD 75
- Disabling the Password terminal 122
- Disk Access Delay Time 139
- Disk Array Parameters 159
- Disk Reserved Space 39, 96
- Disk Reserved Space LCD 53
- Download Boot Record and Firmware 198
- Download FW/BR 197, 198
- Drive I/O Timeout 184
- Drive I/O timeout 190

- Drive Identification 189
- Drive Information terminal 129
- drive members 27
- Drive Motor Spin-up 190
- Drive Predictable Failure Mode 181
- Drive Roaming 190
- Drive Space Allocation 108
- Drive Swap Check Time 152
- drives
 - SCSI, viewing and editing 113
- Drive-side Parameters Advanced config. 139
- Dynamic cache memory allocation 193
- Dynamic on-lining 195

- E**
- Enclosure Devices 149
- enclosure implementation 149
- Enclosure Monitoring 143
- Environment management 193
- event logs
 - viewing and editing 46, 47
- Event Logs terminal 89
- event message 14
- Event Trigger Operations 145
- Event Trigger Option 145, 146
- Event Triggered Operations 143
- Execution Period 155
- Expand Logical Drive 168
- Expand logical drives 166
- Expand Logical Volume 170
- expansion interrupted 165
- Expansion Re-Striping 162

- F**
- Failback 193
- Failed Drive Swap Detect 17
- Failover 193
- Fan failure 143
- Fault Management 29
- Fault Prevention 176
- firmware
 - upgrading 197
- Firmware Binaries 197
- Firmware Binary 198
- Firmware Functionality 186
- Firmware Sync-version 197
- flash memory 197
- flashing
 - all but selected drives 116
 - all SCSI drives 115
 - selected SCSI drives 115

- G**
- Gateway 40
- Gauge Range 82
- Global Spare Drive 187
- global spare drive 15
- global spare drive deleting 114
- Global Spare LCD 61
- Global Spare terminal 106
- GMT Greenwich Mean Time 128
- Guaranteed Latency I/O 190

H

- hard drives 96
- Head and Sector counts 137
- Head Ranges 138
- head sector count 137
- hex number 127
- host application 23, 29
- Host LUN Geometry 190
- Host-side Maximum Queued I/O count 190
- How does it work?
 - the RAID controller 131

I

- I/O timeout
 - SCSI 140
- ID
 - SCSI, deleting 119
- Identifying Drives 115
- identifying drives 115
- Immediate array availability 187
- In-band SCSI 134, 197
- in-band SCSI 134
- Initial Screen 82
 - terminal 82
- initialization 86
- Initialization Mode 56, 97
 - LCD 54
- Internet Protocol 39, 40
- Interruption to RAID expansion 162
- IO channel diagnostics 190
- IP address 197, 198
- IP Address assignment 39
- Iteration Count 67, 111

J

- JBOD 11
 - single drive control 11

L

- LAN port 197
- LBA 111
- LCD Title Display Controller Name 78
- LCD Title Display Controller Name LCD 78
- LCD Title Display Controller Name terminal 125
- LD Initialization Writes 159
- LD Rebuild Writes 159
- Limitations on array config 23
- Limitations on optimization modes 48
- Local Spare 83
- local spare drive deleting 114
- Local Spare LCD 61
- Local Spare terminal 105
- logical drive
 - assigning a name 65
 - deleting 63
 - expansion in Windows 2000® 170
 - explained 9
 - rebuilding 65
- logical drive assigning a name 109
- logical drive deleting 108
- Logical Drive explained 9
- Logical drive identification 186
- logical drive parity check 66

- Logical Drive Parity Regeneration 158
- Logical Drive Preferences 52
 - logical drive preferences 54
- Logical Drive Preferences LCD 52
- Logical Drive Preferences terminal 96
- Logical Drive Status 84
 - terminal 84
- logical drive viewing drive members 63, 107
- logical volume 9, 19, 20, 21, 22
 - explained 9
- Logical Volume Status 44
 - LCD 44
- Logical Volume Status terminal 85
- Logical Volumes 25
- Low-level format 189
- LUN
 - explained 26, 131
 - mappings, viewing and deleting 112
- LUN Applicability 134
- LUN Application 134
- LUN entries 104
- LUN-0's only 134
- LUNs per host ID 133
- LUNs per Host SCSI ID 133

M

- MAC addresses 127
- Main Menu 83
 - terminal 83
- management
 - drive failure 15
- Manual Rebuild 18
- manual rebuild 18
- Mapping a Logical Volume to Host LUNs terminal 103
- Mapping a Logical Volume/Logical Drive to Host LUN LCD 59
- Mapping Partitions 28
- Maximum Array Capacity 186
- Maximum array Number 186
- Maximum Concurrent Host LUN Connection 131
- Maximum concurrent Host LUN connection 190
- Maximum Drive Capacity 53
- Maximum Drive Capacity LCD 53
- Maximum Drive Capacity terminal 96
- maximum drive expand capacity 168
- Maximum Drive Response Time 184, 190
- maximum free capacity 110
- maximum latency 184
- Maximum number of logical volumes 186
- Maximum number of LUNs 186
- Maximum number of LUNs per Host ID 186
- Maximum number of partitions 186
- Maximum Queued I/O Count 133
- Maximum Tag Count (Tag Command Queuing 141
- Media Scan 67, 111, 189
- Media Scan Priority 111
- Mode-1 RAID Expansion 192
- Mode-2 RAID Expansion 192
- Monitoring Thresholds 147
- Multi-level RAID 22
- multiple drive failures 18
- Multiple IDs 103

Multiple target IDs 119
Multi-Threaded Operation 191
mute

beeper 74

Mute Beeper 74

mute beeper 121

Mute Beeper LCD 74

Mute Beeper terminal 121

N

naming controller 125

naming logical drive 65

Native Command Queuing 186

NetMask 40

NEW DRV 86

Normal Drive Writes 159

Notification 200

Number of Tags Reserved for each Host-LUN Connection 132

NVRAM 189

O

Off-Line Mode Initialization 54

On-line Expansion 192

online initialization 98

On-Line Mode Initialization 54

Operational Theory 26

Optimization for Random or Sequential I/O 92

Optimization Mode 24, 30, 48, 91

Optimization Mode LCD 48

Out-of-Band 38, 39

Overwrite Inconsistent Parity 158

P

parameters

cache 48, 91

parameters cache 48

parameters controller 125

parity blocks 110

parity check logical drive 66

parity data 184

Parity Error Event 158

Parity regeneration 158

Parity Synchronization 193

partition

logical drive, deleting 64

partition logical drive, deleting 108

Partitioning a Logical Drive/Logical Volume LCD 58

Partitioning a Logical Drive/Logical Volume terminal 101

partitioning the logical volume 21

Partitions the number of 25

password

changing 74

password changing 121

password disabling 75, 122

password setting a new 122

Password Validation Timeout 78

Password Validation Timeout LCD 78

Password Validation Timeout terminal 127

Periodic Auto-Detect 152

Periodic Cache Flush 93, 191

Periodic Drive Check 182

Periodic Drive Check Time 109, 152, 181

peripheral device 135

Peripheral Device Status 149

Peripheral Device Type 134, 135

Perpetual Clone Clone Failing Drive 177

Perpetual Clone SMART 179

physical drives 23

PID 87

Power supply failure 143

Predictable Failure Test 182

priorities 184

PROCOMM 198

Q

Quick Installation 42

LCD 42

terminal 83

R

RAID 9, 96

advantages 9

level (0+1), disk striping with mirroring 13

level 3, disk striping with dedicated parity disk 14

level 5, striping with interspersed parity 14

RAID (0+1) 18, 22

RAID (3+0) 22

RAID (5+0) 22

RAID 3 14

RAID 5 14

RAID Expansion 161

RAID Expansion Example 170

RAID expansion with logical volume 22

RAID Level 52

RAID Level LCD 52

RAID Levels 10, 186

RAID levels definition of 10

RAID Migration 55

RAID-to-host-bus 197, 198

RAIDWatch Manager 197

Read/Write Test terminal 129

Read/Write testing 189

Read-ahead 191

Rebuild 86

rebuild

logical drive 65

rebuild concurrent in RAID (0+1) 18

rebuild degrade mode 111

Rebuild Priority 159

rebuild priority 62

Rebuild Settings 62

Rebuild Settings LCD 61

rebuild with spare 17

Rebuilding a Logical Drive 65

Rebuilding a Logical Drive LCD 65

Rebuilding Logical Drive terminal 109

Regenerate Parity 194

Regenerating Logical Drive Parity 66

Regenerating Logical Drive Parity LCD 66

Regenerating Logical Drive Parity terminal 110

Replace After Clone 176

Replace after Clone Clone Failing Drive 176

reserved space 130

Reset Controller 75

reset controller 75, 123

Reset Controller LCD 75

- Reset Controller terminal 123
- Restore NVRAM from Disks 76
- Restore NVRAM from Disks LCD 76
- Restore NVRAM from Disks terminal 126
- Restoring the Default Setting 72
- Restoring the Default Setting LCD 72
- RS-232 33
 - configuration via front panel 33
- RS-232 upgrading firmware via 198
- S**
- S.M.A.R.T. 176, 179, 182, 192, 194
- safety range 147
- safety range threshold 148
- SAF-TE 149
 - periodic check time 149
- SAF-TE and S.E.S. Enclosure Monitoring 149
- SAF-TE/S.E.S. polling period 195
- Saving NVRAM to Disks 76
- Saving NVRAM to Disks LCD 76
- Saving NVRAM to Disks terminal 126
- Scan Mode 156
- Scan Priority 156
- Scatter / Gather 191
- Scheduled Maintenance 154
- Scheduled Media Scan 194
- SCSI
 - channel ID, deleting 71
 - channels, viewing and editing 118
 - drives, viewing and editing 113
 - ID, explained 131
- SCSI channel explained 131
- SCSI Channel's Status terminal 87
- SCSI Drive Read/Write Test LCD 80
- SCSI Drive Status 44
 - LCD 44
- SCSI Drive Utilities 80
- SCSI Drive Utilities LCD 80
- SCSI Drive Utilities terminal 129
- SCSI Drive's Status terminal 86
- SCSI ID, deleting 119
- SCSI IO Timeout 140
- SCSI Parameters 30
- Sector Ranges 137
- select "Controller Peripheral Device Configuration 147
- sequential I/O 93
- Serial Port 33
- serial port
 - connection and setup 33, 39
- SES 149
- Set Peripheral Device Entry 145, 146
- Setting a New Password terminal 122
- Shutdown Controller 75
- Shutdown Controller LCD 75
- Shutdown Controller terminal 123
- SMART 176, 179, 180, 181, 182, 192
- SMART attributes 179
- source drive 167, 176, 177, 183
- Spare Drive 51
- Spare drive 21, 86
- spare drive 24, 183
- Spare Drive Assignments 53

- Spare Drive Assignments LCD 53
- Spare Drives 15, 18
- Spin-Up Parameters 31
- SSH 41, 196
- SSL 195
- STAND-BY 86
- stand-by 176, 178
- standby 105
- standby spare 176
- Start Time and Date 155
- streaming 185
- Stripe Size 97, 191
- stripe size 91, 186
- Stripe Size for different RAID levels 49
- Stripe Size LCD 54
- Stripe Size Preset Values 91
- striping
 - with interspersed parity 14
 - with mirroring 13
 - with parity disk 14
- Sun Solaris configuration 137
- Support Removable media 134
- Synchronized cache 193
- System Functions 74, 198
- System Functions LCD 74
- System Functions terminal 121
- T**
- Tag Command Queue 190
- Tag Command Queuing 141, 186
- Tags Reserved 190
- target member drive 178
- Task Schedule 157
- Task Scheduler 154
- Task Scheduler Options 155
- Telix 198
- Temperature exceeds threshold 143
- terminal emulation
 - terminology 83
- the "Password Validation Timeout 74
- the Destination Drive 167
- Time Zone 128
- timeout 184
- Timeout default 140
- transfer protocol 198
- Transfer Rate 82
- Transfer Rate Indicator 82
- Transparent reset 189
- Trigger Events 189
- Trigger Thresholds 148
- Tunable Parameters 29
- U**
- upgrading
 - firmware 197
- upgrading using RAIDWatch 197
- used drive 176
- USED DRV 86
- User Configurable Geometry range 190
- V**
- validation timeout 122
- Verification on Writes 159, 194
- View and Edit Event Logs 47

LCD 47

view clone progress 177

View Connected Drives 51

LCD 51

View Peripheral Device Status 147

Viewing and Editing Host LUNs LCD 69

Viewing and Editing Host LUNs terminal 112

Viewing and Editing Logical Drive and Drive Members terminal 107

Viewing and Editing Logical Drives LCD 63

Viewing and Editing SCSI Channels 70

viewing and editing SCSI channels 118

Viewing and Editing SCSI Channels LCD 70

Viewing and Editing SCSI Channels terminal 118

Viewing and Editing SCSI Drives 69

Viewing and Editing SCSI Drives LCD 69

Viewing and Editing SCSI Drives terminal 113

Viewing and Editing SCSI IDs terminal 118

Viewing or Deleting LUN Mappings terminal 112

Viewing the Connected Drives terminal 94

Voltage and Temperature Parameters 147

volume capacity 101

VT-100

connection and setup 33, 39

terminal 33

VT-100 Mode 82

W

Warning 200

Write Policy 56, 68, 82, 111

write policy 186

Write Policy array 96

Write Policy LCD 53

Write-Back 49

Write-Back cache 68

Write-Back/Write-Through Cache 92

Write-Through 49

write-through caching 184

Z

ZMODEM 198, 199

Bosch Security Systems

Robert-Koch-Straße 100

D-85521 Ottobrunn

Germany

Telefon (089) 6290-0

Fax (089) 6290-1020

www.bosch-securitysystems.com

© Bosch Security Systems, 2006